

VECTOR LOGIC

AI & SOFTWARE DEVELOPMENT — ODISHA, INDIA

PREPARED FOR PAPPARTI

Papparti — Report Appendices

DATE

11 October 2026

CONFIDENTIALITY

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

Contents

Papparti — Report Appendices	3
E. Appendix — every figure and where it came from	3
E.1 One-time build rows	3
E.2 Monthly rows, after handover and go-live	5
E.3 Key rates and inputs	6
F. Appendix — what is measured today and what is not	8
F.1 MEASURED — a real instrument produced these numbers	8
F.2 NOT MEASURED — no instrument has produced these numbers yet	9
G. Appendix — open items and what we need from you	10
H. Closing page	13

Papparti — Report Appendices

Project: Papparti — on-demand any-service marketplace with a self-hosted voice and avatar layer **Prepared by:** Vector Logic — AI & software development, Odisha, India **Prepared for:** Papparti (client) **Date:** 11 October 2026 (IST) **Currency:** INR, Indian numbering. USD converted at the pinned rate 1 USD = ₹96.7384. **Status:** Appendices to the client package. Each figure names the document and section it comes from.

E. APPENDIX — EVERY FIGURE AND WHERE IT CAME FROM

Status key: **MEASURED** — a real instrument produced the number. **ESTIMATED** — a modelled figure with its assumption stated. **UNVERIFIED** — not confirmed; confirm before relying on it. Figures marked LIVE were read from a vendor's own pricing page on 11 October 2026 and are re-priced if the quote expires.

E.1 One-time build rows

#	FIGURE	AMOUNT (₹)	WHERE IT COMES FROM	STATUS
1	Phase 0 — feasibility benchmark	4,47,967	Quotation \$A row 1; Estimate \$1a row 1 / \$3; Proposal \$5	ESTIMATED (GPU live, engineering estimate)
2	Model work — four workstreams	30,90,471	Quotation \$A row 2; Estimate \$1a row 2 / \$4; Proposal \$5	ESTIMATED
3	Software — Phase 1	16,25,205	Quotation \$A row 3; Estimate \$1a row 3 / \$3 (12 weeks)	ESTIMATED
4	Software — Phase 2	10,83,470	Quotation \$A row 4; Estimate \$1a row 4 / \$3 (8 weeks)	ESTIMATED
5	Software — Phase 3	10,83,470	Quotation \$A row 5; Estimate \$1a row 5 / \$3 (8 weeks)	ESTIMATED
6	Deployment + handover	5,41,735	Quotation \$A row 6; Estimate \$1a row 6 / \$3 (4 weeks)	ESTIMATED

#	FIGURE	AMOUNT (₹)	WHERE IT COMES FROM	STATUS
7	One-time subtotal, excluding GST	78,72,318	Quotation §A subtotal; Estimate §1a; Proposal §6.1; Commercial Terms §2/§4.1	MEASURED (sum of the rows)
8	GST @ 18%	14,17,017	Quotation §A; Proposal §6.1; Commercial Terms §4.1 (= 78,72,318 × 0.18)	MEASURED (arithmetic)
9	One-time total, including GST	92,89,335	Quotation §A; Estimate §1a; Proposal §6.1; Commercial Terms §4.1	MEASURED (arithmetic)
10	Phase 4 — real-time lip-synced avatar (option)	13,54,338	Quotation §B row 7; Estimate §3 / §10; Proposal §5, §6.4	ESTIMATED ("starting estimate, not a quote")
11	Model work — STT workstream	8,14,886	Estimate §4	ESTIMATED
12	Model work — TTS workstream	8,13,973	Estimate §4	ESTIMATED
13	Model work — Reasoning LLM workstream	10,93,744	Estimate §4	ESTIMATED
14	Model work — Voice production workstream	3,67,868	Estimate §4; Proposal §3.4; Commercial Terms §8	ESTIMATED
15	Voice production — engineering (2 work-weeks)	2,70,868	Proposal §3.4; Estimate §4	ESTIMATED
16	Studio recording — ~20 h delivered audio	97,000	Proposal §3.4; Estimate §4 / §10; Commercial Terms §2.4, §8	ESTIMATED (midpoint of a ₹64,000–₹1,30,000 band)
17	Model work — GPU cost subtotal (244 GPU-hours)	13,927	Estimate §4	ESTIMATED (hours) / LIVE (rate)
18	Model work — engineering subtotal	29,79,544	Estimate §4	ESTIMATED

#	FIGURE	AMOUNT (₹)	WHERE IT COMES FROM	STATUS
19	Software development share of one-time total (Phases 1-3)	37,92,145	Estimate §8; Proposal §6.6	MEASURED (arithmetic, 51% share)
20	Mobilisation advance (10% of one-time fee)	7,87,232	Commercial Terms §2.1, §2.4	MEASURED (arithmetic)
21	Milestone 3 net (Model work less advance recovered)	23,03,239	Commercial Terms §2.1	MEASURED (arithmetic)
22	Retention / holdback (5% of one-time fee)	3,93,616	Commercial Terms §2.3	MEASURED (arithmetic)
23	Handover invoice net of retention	1,48,119	Commercial Terms §2.3 (5,41,735 – 3,93,616)	MEASURED (arithmetic)
24	Phase 0 — GPU rental sub-charge	41,665	Estimate §3; Proposal §4	LIVE (RunPod @ \$0.59/hr × 730 hr)
25	Phase 0 — engineering sub-charge	4,06,301	Estimate §3; Proposal §4 (3 work-weeks)	ESTIMATED
26	Change-order absorption threshold	13,542	Commercial Terms §5.2 (4 engineering hours)	MEASURED (arithmetic)

E.2 Monthly rows, after handover and go-live

#	FIGURE	AMOUNT (₹/MONTH)	WHERE IT COMES FROM	STATUS
27	GPU serving — 1× 48 GB dedicated, 24/7	41,665	Estimate §1b / §5; Quotation §C row 1; Commercial Terms §6	LIVE (RunPod @ \$0.59/hr × 730 hr)
28	Managed PostgreSQL + vector store	5,891	Estimate §1b / §5; Quotation §C row 2; Commercial Terms §6	LIVE (DigitalOcean @ \$60.90/mo)
29	Redis	1,935	Estimate §1b / §5; Quotation §C row 3; Commercial Terms §6	LIVE (Upstash @ \$20/mo)

#	FIGURE	AMOUNT (₹/MONTH)	WHERE IT COMES FROM	STATUS
30	Real-time media routing (SFU/TURN)	4,837	Estimate §1b / §5; Quotation §C row 4; Commercial Terms §6	LIVE (managed real-time media tier @ \$50/mo)
31	Object storage	1,161	Estimate §1b / §5; Quotation §C row 5; Commercial Terms §6	ESTIMATED (~500 GB at a live rate; volume assumed)
32	Monitoring + CI	2,902	Estimate §1b / §5; Quotation §C row 6; Commercial Terms §6	ESTIMATED (~\$30/mo placeholder)
33	Support retainer (0.5 work-week / month)	67,717	Estimate §1b / §5; Quotation §C row 7; Commercial Terms §6	ESTIMATED (assumption at the engineering rate)
34	Monthly total	1,26,108	Quotation §C; Estimate §1b; Proposal §6.2; Commercial Terms §6	MEASURED (sum of the rows)
35	Non-GPU, non-retainer monthly lines	16,726	Commercial Terms §6.1 (5,891 + 1,935 + 4,837 + 1,161 + 2,902)	MEASURED (arithmetic)
36	Second 48 GB serving GPU, if concurrency requires	~83,000	Estimate §5; Proposal §6.2; Quotation §D (not included)	UNVERIFIED (flagged, not priced)

E.3 Key rates and inputs

#	FIGURE	AMOUNT	WHERE IT COMES FROM	STATUS
37	FX rate — USD to INR	1 USD = ₹96.7384	Estimate header / §6 / §9 [5]; Proposal §10; GPU Decision §8 (A1)	LIVE (XE mid-market, 11 Oct 2026, 04:05 UTC; re-verified 96.7372)
38	Engineering rate	\$35/hour	Estimate §6; Proposal §6.0 / §10; Commercial Terms §5.1	ESTIMATED (internal assumption; CEO must confirm)
39	Engineering rate — hourly in INR	₹3,385.84/hour	Estimate §6; Commercial Terms §5.1	MEASURED (arithmetic on the pinned FX)

#	FIGURE	AMOUNT	WHERE IT COMES FROM	STATUS
40	Engineering rate — per day	₹27,087/day	Estimate \$6	MEASURED (arithmetic)
41	Engineering rate — per work-week	₹1,35,434 per 40-hour week	Estimate \$6 / \$9; Proposal \$10; Quotation \$F; GPU Decision \$6	MEASURED (arithmetic)
42	Work-week	40 hours, 5 days	Estimate \$9; Proposal \$10	MEASURED (standard)
43	GPU hours per month	730 (24×7)	Estimate \$9; Proposal \$10; GPU Decision \$1.1	MEASURED (24 × 7 × 365/12)
44	GPU hourly rate — reference serving	\$0.59/hour (48 GB)	Estimate \$2 [1]; GPU Decision \$8 (A2)	LIVE (RunPod, 11 Oct 2026)
45	GPU hourly rate — Vast.ai floor / median	\$0.28 / \$0.40 per hour	Estimate \$2 [2]	LIVE (not the planning rate)
46	GPU hourly rate — Lambda Labs	\$0.80 per hour	Estimate \$2 [3]	LIVE (not the planning rate)
47	GPU hourly rate — India-hosted tier	\$1.20 / \$1.44 per hour	Estimate \$2 [4] (L40S / A40 48 GB)	LIVE (contingency, not default)
48	Stale GPU figure corrected out	₹34,056/month (at ₹88/USD)	Quotation \$F; Estimate \$2; Commercial Terms \$1	UNVERIFIED (superseded; kept only to show the correction)
49	Payment-gateway fee	2% + 18% GST on the fee	Estimate \$7b; Proposal \$6.4; Commercial Terms (via \$6 note)	LIVE (Razorpay structure)
50	Gateway fee worked example on ₹1,00,000 settled	₹2,360	Estimate \$7b; Proposal \$6.4	MEASURED (arithmetic)
51	Commission floor the gateway implies	~2.4% of booking value	Estimate \$8; Proposal \$8.3	ESTIMATED (arithmetic on the live fee)
52	GPU break-even — used card, refurb server (Option C)	~21 months	Quotation \$F; GPU Decision \$1.3 / \$4	ESTIMATED (modelled)

#	FIGURE	AMOUNT	WHERE IT COMES FROM	STATUS
53	GPU break-even — used card, India OEM server (Option B)	~34 months	GPU Decision \$1.3 / \$4	ESTIMATED (modelled)
54	GPU break-even — new card, new server (Option A)	~61 months	Quotation \$F; GPU Decision \$1.3 / \$4	ESTIMATED (modelled)
55	Rent vs buy — 36-month total, renting	₹14,99,940	GPU Decision \$1.3	ESTIMATED (modelled)
56	Rent vs buy — 36-month total, Option C	₹10,03,217	GPU Decision \$1.3	ESTIMATED (modelled)
57	Odisha LT-industrial energy charge	₹6.20/kWh	GPU Decision \$2e / \$8 (A4)	LIVE (OERC RST FY 2026-27)

Every row above names a real document and section. The totals in E.1 and E.2 stand independently of the UNVERIFIED and ESTIMATED rows: where a figure is an assumption it is not added to a verified sum — the arithmetic is shown where it is used.

F. APPENDIX — WHAT IS MEASURED TODAY AND WHAT IS NOT

F.1 MEASURED — a real instrument produced these numbers

1. **Voice sample duration — male greeting** ~5.5 s (Proposal §8.4).
2. **Voice sample duration — female greeting** ~6.8 s (Proposal §8.4).
3. **Voice sample duration — male number sentence** ~6 s, containing a price and a duration (Proposal §8.4).
4. **Voice sample duration — female number sentence** ~6 s (Proposal §8.4).
5. **Male voice speech energy on the number sentence** -16.2 dBFS — further from clipping (Proposal §8.4).
6. **Female voice speech energy on the number sentence** -12.8 dBFS (Proposal §8.4).
7. **Male voice speech energy on prose samples** -15.4 dBFS (Proposal §8.4).
8. **Proportion of non-silent frames**, measured per voice sample (Proposal §8.1 — audio generated and measured objectively).

9. **Real-time factor (RTF)** of the voice output on our own hardware (Proposal §8.1 — measured).
10. **Avatar tonal steps — high tier** 30 distinct steps spanning 131 of 255 (Proposal §8.3).
11. **Avatar tonal steps — low tier** 1 distinct step spanning 0 (Proposal §8.3) — a measured $30\times$ difference, enforced by a script that fails the build if it collapses.
12. **Avatar mouth-region redraw between idle and talking frames** 35.8% (Proposal §8.3).
13. **Horizontal overflow** 0 px at 1920 / 1440 / 1024 / 390 px widths (Proposal §8.3).
14. **GPU rental price read from the vendor's live page** \$0.59/hour for a 48 GB card on 11 October 2026 (Estimate §2; GPU Decision §1).
15. **FX rate read live** 1 USD = ₹96.7384, cross-checked at 96.7372 on 11 October 2026 (Estimate §9; GPU Decision §8).
16. **All totals and their arithmetic** — the one-time subtotal, the GST line and inclusive total, the monthly total, the payment-schedule sum, and the milestone figures were all recomputed and pass (Commercial Terms §11; Estimate §8).
17. **The renders and demo artefacts shipped with the proposal** — four avatar images, the speaking video, and the scripted conversation clips all exist as files (Proposal §6.0).

F.2 NOT MEASURED — no instrument has produced these numbers yet

1. **The Phase 0B benchmark as a whole has not run.** It is the gating measurement of the entire plan and is the next paid step, not a completed result (Proposal §4; Commercial Terms §3.1).
2. **Hindi recognition accuracy (best base, after fine-tune)** — not measured; the gate is 12% WER (Proposal §4).
3. **Hinglish code-switch recognition accuracy (before any fine-tune)** — not measured; the gate is 20% WER (Proposal §4).
4. **Fine-tune improvement over the untrained base** — not measured; the gate is $\geq 15\%$ relative gain (Proposal §4).
5. **Sessions per node under real load** — not measured; the gate is ≥ 10 . A $3\times$ error here triples the infrastructure bill (Proposal §4; GPU Decision §5).
6. **p95 time to first audio** — not measured; the gate is $\leq 1,500$ ms (Proposal §4).
7. **Booking tool-call exact-match accuracy (Hindi/Hinglish)** — not measured; the gate is 85% (Proposal §4).

8. **Ambiguous-duration inference rate** — not measured; the gate is a hard stop above 0% (Proposal §4).
9. **General-capability regression after fine-tuning** — not measured; the gate is under 3% (Proposal §4).
10. **No MOS and no listening test exists for the voice output.** No human listener has rated it, and no quality score is quoted (Proposal §8.1; Commercial Terms §9.1).
11. **The client's real GPU utilisation** — unknown until Phase 0B measures it. This is why renting is recommended rather than buying (GPU Decision §5, §7).
12. **Concurrency capacity of a single 48 GB card under real load** — unmeasured; the monthly GPU line assumes single-stream serving (Estimate §11; GPU Decision §5).
13. **Any Hindi speech quality number from a model card** — every figure in our research is a vendor claim or a model-card number, not a measurement (Proposal §10; Estimate §11).
14. **The two conflicting Hindi accuracy figures in our research** — cannot be reconciled from the documentation; only our own measurement replaces them (Proposal §8.1).
15. **Vendor-specific latency for India-served inference** — not measured; the India-hosted GPU stays a contingency (Estimate §11).
16. **No human judgement of the avatar's art direction** — that decision belongs to the client (Proposal §8.3).
17. **The studio recording has not run** — the final brand voice does not exist yet (Proposal §6.0).
18. **Node power draw (600 W) and the real GPU-colo price** — assumptions, not measurements (GPU Decision §7).

G. APPENDIX — OPEN ITEMS AND WHAT WE NEED FROM YOU

Every field below is blank on purpose. Each is a fact only a human can supply; none may be filled by guesswork. The documents are not signature-ready until these are complete.

#	OPEN ITEM	WHO MUST SUPPLY IT	WHY IT CANNOT BE GUESSED
1	Client legal entity name as registered	Papparti	The address reads "Papparti", a brand. A contract and every invoice need the registered legal person.
2	Signatory name and designation	Papparti	Milestone acceptance must be unambiguous, and the right person must receive each deliverable.
3	Contact email (the address that reaches you)	Papparti	contact@papparti.com is used only as a placeholder in the demo references.
4	Client registered address (for the tax invoice)	Papparti	Required on a valid Indian tax invoice; not supplied.
5	Client GSTIN (if registered)	Papparti	A registration number must be read from the certificate, never constructed.
6	Client PAN (for TDS matching)	Papparti	A number to be read from the client's records.
7	Vector Logic GSTIN	Vector Logic	To be read from the registered entity's GST certificate before the first invoice.
8	Vector Logic PAN	Vector Logic	To be read from Vector Logic's records, so TDS credit can be matched.
9	SAC code — software development service (cit. 998314)	Tax adviser (both parties)	Recorded UNVERIFIED ; the current code must be confirmed before it prints on an invoice.
10	SAC code — AI / model training services	Tax adviser (both parties)	No code is asserted at all; the field is blank on purpose.
11	Place of supply and CGST/SGST vs IGST split	Tax adviser	Determined under the IGST Act per invoice; not asserted as a settled position.

#	OPEN ITEM	WHO MUST SUPPLY IT	WHY IT CANNOT BE GUESSED
12	TDS section and rate	Tax adviser (both parties)	The applicable section and rate must be confirmed under Indian law.
13	Voice artist identity and release date	Papparti (with counsel)	A release is signed by a named artist; no release = no clone.
14	Does real Hindi/Hinglish audio already exist? (call recordings, voice notes)	Papparti	Changes Phase 0's test design; if there is none, that is workable but must be known.
15	The service taxonomy — the actual launch service categories	Papparti	Drives the recognition vocabulary and the reasoning model's structured output.
16	The locality names — real place-names and their Hindi pronunciations	Papparti	Drives the recognition vocabulary and the locality field.
17	One launch vertical and one launch locality	Papparti	"Any service" from day one is the measured failure mode of this sector.
18	Demo script's city vs your market	Papparti	The demo references a city; confirm it matches the locality you intend to launch.
19	Payment-days convention (e.g. Net 15 / Net 30)	Both parties	A commercial convention the parties agree; not assumed.
20	Late-payment interest rate	Both parties	Agreed in the contract; not assumed.
21	Acceptance-review window (days, for deemed acceptance)	Both parties	Agreed in the contract; not assumed.
22	Governing-law city / jurisdiction	Both parties	The seat of the contract is a negotiated choice.
23	Effective date of the contract	Both parties	The date of execution.
24	30-day warranty start date	Both parties	Fixed by the handover acceptance certificate.

Items 1–6 and 13–18 are the ones that block the earliest work. Items 7–12 and 19–24 block signature and invoicing.

H. CLOSING PAGE

If you want to proceed. Two paths, and nothing else is needed to start.

Path one — proceed as scoped. Supply your registered legal entity name and your signatory's name and designation, and confirm your acceptance in writing on or before **10 November 2026**. Vector Logic adds its GSTIN and the confirmed SAC codes at invoice time. Work begins with Phase 0: three work-weeks of engineering plus one month of GPU time, delivering a written feasibility report with our own measured numbers and a go/no-go on each of the three models. Phase 1 does not begin until Phase 0's results are accepted. The mobilisation advance of ₹7,87,232 is drawn against setup and is recovered at the Model work milestone.

Path two — change something first. Raise it in writing before acceptance. A change to the scope is described as a discrete deliverable and priced before any work begins, at the shown rate of ₹3,385.84/hour (₹1,35,434 per 40-hour work-week), plus any pass-through vendor cost at the then-current live price, plus GST. Changes whose total value is at or below ₹13,542 are absorbed at no charge. If the change is a scope-down to meet a lower budget, the model workstream is the line to protect, because it is the part no one else is selling you.

Either way, one date governs: this pricing is valid until **10 November 2026**, because the monthly figures rest on live vendor and exchange rates read on 11 October 2026. After that date Vector Logic re-prices the affected lines from then-current prices and issues a fresh validity window. The engineering rate is the part that does not move.

Prepared by Vector Logic, 11 October 2026. All figures traceable to the client package. Base model identities are commercially confidential to Vector Logic and are not disclosed in this document.