

PREPARED FOR PAPPARTI

Papparti — Commercial Terms

(Annexure to the Project Proposal)

The commercial annexure: what gets paid, when, and what makes a milestone accepted.

For legal and finance review: milestones, acceptance, tax, termination, voice rights.

Asks for: the blank fields listed in the open-items register.

AT A GLANCE

Payment	Six milestones, each tied to an accepted artefact
Advance	10 percent, recovered against the model-work milestone
Retention	5 percent, released at handover
Phase 0 acceptance	On measured benchmark results, not on opinion
Blank fields	Left blank on purpose, never guessed

DATE 11 October 2026	VALIDITY 30 days; expires 10 November 2026.	CONFIDENTIALITY Confidential. Prepared by Vector Logic for Papparti. Not for distribution beyond Papparti and its advisers.
--------------------------------	---	---

Contents

Papparti — Commercial Terms (Annexure to the Project Proposal)	4
1. Validity	4
2. Payment schedule	4
2.1 The schedule	4
2.2 How the advance works	5
2.3 Retention / holdback	5
2.4 Why a mobilisation advance, and where in the industry range we sit	5
2.5 Milestones are tied to accepted artefacts, not time	6
2.6 GST on each invoice	6
2.7 Late payment	6
3. Acceptance criteria per milestone	6
3.1 Milestone 2 — Phase 0 (feasibility benchmark)	6
3.2 Milestone 3 — Model work (four workstreams)	7
3.3 Milestone 4 — Phase 1 (voice core + booking engine)	7
3.4 Milestone 5 — Phase 2 (avatar + polish)	7
3.5 Milestone 6 — Phase 3 (second vertical / locality)	7
3.6 Milestone 7 — Deployment + handover	8
3.7 Handover acceptance certificate	8
3.8 Acceptance deemed given	8
3.9 What happens if a Phase 0 gate FAILS — the critical clause	8
4. Invoicing and tax	9
4.1 The Figures Invoiced	9
4.2 GSTIN	9
4.3 Place of supply	9
4.4 HSN / SAC codes — UNVERIFIED	9
4.5 TDS	10
4.6 Invoice mechanics Vector Logic requires	10
5. Scope-change mechanism	10
5.1 How a change is raised and priced	10
5.2 The absorption threshold	10
5.3 What is not a change	11
6. What triggers the monthly charge	11
6.1 The gap between handover and the first paying customer	11
7. Termination	12
7.1 If the client stops	12
7.2 If Vector Logic cannot deliver	12
7.3 Trained model artefacts on termination	12
7.4 Termination for convenience	13
8. The voice-rights release requirement	13
8.1 What Vector Logic does if no release is provided	13
9. Liability and warranties	14
9.1 The accuracy limitations disclosed in the proposal survive into the contract	14
9.2 What Vector Logic does warrant	14
9.3 What Vector Logic does not warrant	14
9.4 Bounded liability	15

10. Open fields register	15
11. Verification	16

Papparti — Commercial Terms (Annexure to the Project Proposal)

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) **Status:** Annexure to CLIENT-PROPOSAL.md. To be read together with it and with CLIENT-ESTIMATE.md.

This annexure is a commercial document, not legal advice. It records the commercial terms Vector Logic proposes, in plain language, so that both parties and their advisers can review the same numbers. Every clause below is subject to the written contract that follows signature, and the parties should each have their own counsel and chartered accountant review it before signing. Nothing in this document creates a legal obligation on its own.

Every figure in this annexure is traceable to CLIENT-ESTIMATE.md, or is arithmetic on those figures which is shown where it is used. No figure has been invented for this annexure. Base AI model identities are commercially confidential to Vector Logic and are deliberately not named anywhere in this document; models are referred to by capability and parameter class only.

1. VALIDITY

This pricing is valid for 30 days from 11 October 2026 — that is, until 10 November 2026 (inclusive).

The reason is commercial, not procedural: **the monthly operational figures rest on live cloud GPU and managed-hosting prices that were read from the vendors' own pricing pages on 11 October 2026, and those prices move.** The GPU serving line alone (₹41,665/month) is a live vendor rate of \$0.59/hour converted at a live FX rate of ₹96.7384/USD; both the rate and the exchange rate are re-read from the vendor and the market, not fixed by us. Our own cost estimate already had to correct a previously-cited GPU figure that had gone stale (₹34,056 → ₹41,665, roughly 22% higher). A quote built on live prices has a shelf life, and 30 days is it.

What happens after expiry. If the offer has not been accepted in writing within the validity window, Vector Logic will **re-price the quotation at the vendor and FX prices then current**, and issue a fresh validity window. The re-price applies to the live-priced components (the monthly running costs, and the GPU sub-charge inside Phase 0) and, if the FX rate has moved materially, to the INR value of the USD-billed lines.

What does not move. The **engineering rate is the part that does not move.** The \$35/hour rate (₹3,385.84/hour, ₹1,35,434 per 40-hour work-week, at the pinned rate) is Vector Logic's own commercial rate, not a vendor price, and it holds across re-quotes. The one-time build fee is therefore stable in its engineering component; only the pass-through vendor and FX components are re-priced at expiry.

2. PAYMENT SCHEDULE

The one-time project fee is **₹78,72,318 excluding GST**. It is invoiced in **six milestones**, each tied to a **delivered, accepted artefact** — never to elapsed time. A milestone is invoiced when its artefact is delivered and accepted under §3.

2.1 The schedule

#	MILESTONE	MAPS TO PHASE (FROM CLIENT-ESTIMATE.MD)	INVOICE INR (EXCL. GST)
1	Mobilisation advance (10% of the one-time fee)	Drawn against Phase 0 setup, GPU procurement, the studio booking and Phase 0B benchmark preparation	7,87,232

#	MILESTONE	MAPS TO PHASE (FROM CLIENT-ESTIMATE.MD)	INVOICE INR (EXCL. GST)
2	Phase 0 — feasibility benchmark	Phase 0 (₹4,47,967)	4,47,967
3	Model work — four workstreams	Model work (₹30,90,471) less the ₹7,87,232 advance recovered	23,03,239
4	Phase 1 — voice core + booking engine	Phase 1 (₹16,25,205)	16,25,205
5	Phase 2 — avatar + polish	Phase 2 (₹10,83,470)	10,83,470
6	Phase 3 — second vertical / locality	Phase 3 (₹10,83,470)	10,83,470
7	Deployment + handover	Deployment + handover (₹5,41,735)	5,41,735
Sum of all invoices			78,72,318

Sum check: 7,87,232 + 4,47,967 + 23,03,239 + 16,25,205 + 10,83,470 + 10,83,470 + 5,41,735 = **₹78,72,318** — exactly the one-time fee excluding GST. Verified in §11.

Each of the six phase amounts is invoiced at its estimate value; the only adjustment in the schedule is that the **mobilisation advance (₹7,87,232) is recovered against the Model work milestone**, because that is where the front-loaded spend it funded (studio recording, GPU procurement) is actually consumed. The client therefore pays exactly the phase price for each phase's scope, and pays the advance once, not twice.

2.2 How the advance works

The mobilisation advance is **not an extra line and not a fee** — it is an **early draw against the project**, recovered at the Model work milestone. It exists because this project front-loads real cash out of Vector Logic's pocket before there is a deliverable (see §2.4), and it means the client is never paying for the same work twice. The retention (§2.3) is held separately at the other end of the project.

2.3 Retention / holdback

5% of the one-time fee — ₹3,93,616 — is held back as retention and is released at handover and acceptance. In practice the final handover invoice (milestone 7, ₹5,41,735) is settled **less the retention** (₹5,41,735 – ₹3,93,616 = ₹1,48,119 paid at handover), and the retention of **₹3,93,616** is invoiced and released when the client signs the handover acceptance certificate (§3.7), at which point the 30-day warranty has started.

Retention of ₹3,93,616 = 5% × ₹78,72,318, rounded to the rupee. It is the client's protection: it stays in the client's hands until the delivered system is accepted, and it is a real, bounded amount rather than a token. Because the retention is released inside the same milestone-7 line, the total invoiced across §2.1 is unchanged at ₹78,72,318.

2.4 Why a mobilisation advance, and where in the industry range we sit

The advance is **10% of the one-time fee (₹7,87,232)**, and it is justified from the actual scope, not from a template:

- **This project front-loads real cash out of Vector Logic's pocket before there is a deliverable.**

Phase 0 requires a **rented 48 GB GPU for a full month (₹41,665)** to be procured and running before a single line of the benchmark is written. The **studio voice recording (~₹97,000 of delivered audio, a pass-through cost at the midpoint of a ₹64,000–₹1,30,000 modelled band)** has to be booked, cast, directed and recorded before the synthesis model can train at all. And the **Phase 0B benchmark** is a substantial measurement campaign run before the big development spend. None of these exist as a deliverable until after the money is physically spent.

- **The industry range** for mobilisation or advance payments on fixed-scope software and AI

engagements is commonly **10%–30% of contract value up front**, with 10–15% at the conservative end and 20–30% where the supplier has to buy hardware or lengthy data-collection services at the start. **Vector Logic sits at the bottom of that range, at 10%.** The reason we sit low is deliberate: the phase structure already protects Vector Logic (each milestone is invoiced only after an accepted deliverable), so the project does not need a large advance, and the client retains most of its cash against delivered work.

2.5 Milestones are tied to accepted artefacts, not time

No milestone invoice is raised on a date. Each is raised when the named artefact is **delivered and accepted** under §3, or where a gate resolves under §3.9, when the gate resolution itself is accepted.

2.6 GST on each invoice

GST @ 18% is charged on each invoice, and is the client's **input tax credit where applicable** (where Papparti is GST-registered and the service is used for business — see CLIENT - ESTIMATE .md §7c for the reverse-charge mechanics on the imported cloud lines, which are separate from Vector Logic's own GST). GST is not included in any INR figure in §2; each figure is an excluding-GST figure. The GST-inclusive one-time total is in §4.1.

2.7 Late payment

Stated plainly: **each invoice is due within the payment-days convention fixed in the contract (BLANK FIELD — see §10); if an undisputed invoice is not paid by its due date, Vector Logic may pause work on the following milestone until it is paid, and may charge simple interest on the overdue amount at a rate fixed in the contract.** Stated at the same time and with equal weight: Vector Logic will not pause work over a *disputed* invoice — a disputed line is handled under §5 and the remainder of the work continues. The purpose of the late-payment term is not to penalise the client; it is to protect a small agency from carrying a large project on its own balance sheet while vendor GPU bills and studio costs fall due regardless.

3. ACCEPTANCE CRITERIA PER MILESTONE

Every milestone is accepted against a **testable artefact**. Where Phase 0 is concerned, acceptance is on a **measured number** — the recorded Phase 0B results, including pass/fail against the hard gates in MODEL - DESIGN .md §4.4 — and **not on opinion**. The client should not accept Phase 0 on "it looks good"; acceptance is a table of measurements against thresholds.

3.1 Milestone 2 — Phase 0 (feasibility benchmark)

Artefact delivered: a written Phase 0B feasibility report containing, for each model workstream: (i) the frozen test sets used, (ii) the measured metrics, (iii) the recorded **pass/fail against each hard gate**, (iv) the measured **sessions/node** under the concurrency ramp, and (v) a go/no-go on each of the three models.

Accepted when: the report is delivered and it contains a recorded result against **each** of the ten gates below, taken from MODEL - DESIGN .md §4.4:

GATE	THRESHOLD	ON FAIL
Hinglish code-switch WER (STT base, pre-fine-tune)	worse than 20%	Stop/redesign — text-first from day 1
Hindi WER (best base, post-fine-tune)	worse than 12%	Change base
STT fine-tune gain vs stock	under 15% relative	Ship stock + re-scope
TTS MOS (brand voice, ≥30 raters)	below 3.5	Fall back, de-prioritise brand voice

GATE	THRESHOLD	ON FAIL
TTS number/duration intelligibility	below 100%	Hard gate — do not ship synthesis on the money path
LLM tool-call exact-match (Hindi/Hinglish)	below 85%	Change model
Ambiguous-duration inference rate	above 0%	Hard stop on the reasoning model as specified
General-capability regression after fine-tune	above 3%	Retrain with more anchor data
Sessions/node (concurrency, audio-only)	below 10	Re-size the fleet
p95 time-to-first-audio	above 1,500 ms	Re-size / re-design

Acceptance is the **recorded table of measured values**, not a narrative. If a metric was not measured, it is reported as **NOT MEASURED** and the milestone is **not** accepted on that line.

3.2 Milestone 3 — Model work (four workstreams)

Artefact delivered: the three trained adapters (speech recognition, speech synthesis, reasoning) plus the brand-voice master; each adapter packaged with its serving endpoint configuration, and a per-workstream evaluation report measured against the §3.1 gates.

Accepted when: (i) all three adapters load and serve on the target GPU; (ii) each workstream's post-training metrics meet the gate thresholds agreed at Phase 0 acceptance (or the gate resolution in §3.9 was applied); (iii) the **number/duration intelligibility gate shows 100%** and the **ambiguous-duration inference gate shows 0%** — these two are non-negotiable and cannot be waived by a partial credit or an average; and (iv) the voice master is delivered and the client has approved the master take.

3.3 Milestone 4 — Phase 1 (voice core + booking engine)

Artefact delivered: the audio-only voice core (self-hosted speech in and out), the reasoning endpoint, the duration-aware booking engine (30 min / 1 hr / 5 hr), payments with provider split, the immutable ledger, provider onboarding and working hours, and the PWA shell.

Accepted when (the proposal's own Phase 1 gate, made testable): (i) an **end-to-end real booking is made by voice and paid**; (ii) the **provider is paid on completion**; (iii) the **ledger reconciles** to the rupee against the booking and payment records; and (iv) **two overlapping bookings are correctly rejected by the database, not by application logic** — demonstrated by a test that shows the rejection coming from a database constraint, not from a code check that could be removed.

3.4 Milestone 5 — Phase 2 (avatar + polish)

Artefact delivered: the avatar interaction layer on both sides (rendered on the phone), UX polish, and hardening.

Accepted when: (i) the avatar **adds under ~800 ms to p95 latency**; (ii) **captions land inside 500–800 ms**; and (iii) the avatar on the budget-tier rendering path holds its frame budget per the project's own measured rendering numbers. These are measured on the client's own device band, not on a developer workstation.

3.5 Milestone 6 — Phase 3 (second vertical / locality)

Artefact delivered: a second service vertical or locality live on the same platform — catalogue, pricing rules, provider cohort, localisation.

Accepted when: (i) the second vertical or locality is live and can take real bookings end to end; and (ii) **it only starts after vertical 1 is per-order profitable including GPU cost** — the gate the proposal names — with the profitability measured from the ledger, not asserted.

3.6 Milestone 7 — Deployment + handover

Artefact delivered: infrastructure-as-code for the self-hosted stack on rented GPU, CI/CD, monitoring wiring, runbooks, admin training, and the 30-day warranty start.

Accepted when: (i) the stack is stood up from the infrastructure-as-code by the client's own team, following the runbooks, without Vector Logic hands on the keyboard; (ii) monitoring is wired and reaches the client's chosen alerting destination; (iii) the admin console has been handed over with training delivered; and (iv) the code, model weights and endpoints have been transferred.

3.7 Handover acceptance certificate

Handover acceptance is a **signed acceptance certificate** from the client. On signature the **30-day warranty starts** and the **retention of ₹3,93,616 is released** (§2.3).

3.8 Acceptance deemed given

If the client does not respond to a delivered artefact's acceptance request within the acceptance-review window fixed in the contract (BLANK FIELD — §10), the artefact is **deemed accepted** and the milestone is invoiced. This protects Vector Logic from indefinite open milestones; the window is set generously and is a blank field for the client to agree, not a number we have assumed.

3.9 What happens if a Phase 0 gate FAILS — the critical clause

This is the clause that must be right, and it has two sides that both have to hold:

The client must not owe more money for a phase that failed a go/no-go gate. Vector Logic must not be left unpaid for work already done.

It resolves as follows:

1. **Phase 0's deliverable is the measurement itself, not a passing result.** The Phase 0 milestone

(₹4,47,967) is **payable on delivery of the benchmark report with its recorded results**, whether the results pass or fail. The client is buying the answer to the question, and a measured "no" is a complete, valuable answer — it is exactly what Phase 0 exists to produce. Vector Logic is therefore paid for Phase 0 in full, on delivery.

1. **A failed gate stops the dependant spend.** If a gate returns a **STOP/REDESIGN** or **HARD STOP**

result, then **no dependent milestone proceeds** and **no dependent milestone is invoiced**. The client owes **nothing** for Phase 1, Phase 2, Phase 3, the deployment milestone, or the affected part of the model work. The project stops at the Phase 0 price point. The client never pays for downstream work that a go/no-go gate has just invalidated.

1. **A failed gate that changes the model, not the project, is a change order, not a new invoice.**

Gates whose failure means "*change the model*" (Hindi WER > 12%; LLM tool-call < 85%) or "*change the base*" do not end the project — they are resolved by re-scoping under §5. The re-scope is **priced as a change**, and **the client pays nothing extra for the phase that failed the gate**; the change is quoted and approved before any new money is committed.

1. **A HARD STOP** (TTS number/duration intelligibility below 100%, or ambiguous-duration inference

above 0%) **removes the affected capability from the money path**. The synthesis model is held off the pricing path and the booking flow ships with a conventional, on-screen confirmation and receipt — the proposal's own structural safeguard (§2 of

the proposal). The affected work is **re-scoped downward**, and the client is **credited** the difference between the original line and the re-scoped line on the next invoice. The client does not pay for a capability that cannot safely ship.

1. **Work already done that is unaffected by the failure is paid.** If a gate fails in one workstream

but other workstreams (or software already delivered and accepted) are complete and unaffected, those are invoiced and paid normally. The gate failure reduces **only** the dependant scope.

In one sentence: a failed go/no-go gate never adds a rupee to the client's bill and never leaves Vector Logic unpaid for the measurement it delivered; it removes downstream scope and its downstream invoice, and any partial work continues to be paid only where it survives the gate.

4. INVOICING AND TAX

4.1 The Figures Invoiced

ITEM	INR
One-time fee, excluding GST	78,72,318
GST @ 18%	14,17,017
One-time fee, including 18% GST	92,89,335

Arithmetic, shown: ₹78,72,318 × 0.18 = ₹14,17,017 (to the rupee); ₹78,72,318 + ₹14,17,017 = ₹92,89,335. Verified in §11.

Optional Phase 4 (real-time lip-synced avatar) is a separate **₹13,54,338** option, excluded from the one-time total, and its own GST applies on invoicing.

4.2 GSTIN

Every Vector Logic invoice will carry a real GSTIN. The GSTIN is a field to be filled from the registered entity's records before any invoice is issued:

- ****Vector Logic GSTIN:** _____ (BLANK FIELD — must be filled from the registered

entity's GST registration before the first invoice. Not invented here.)**

Client GSTIN (for tax invoice): _____ (BLANK FIELD — to be supplied by Papparti.)

If either party is not GST-registered, that is recorded and the invoicing follows the applicable non-registered treatment, confirmed with a tax adviser.

4.3 Place of supply

Place of supply is determined for each invoice under the IGST Act, and will be recorded on the invoice. As a services supply to a registered recipient, it is normally the recipient's location; the **exact place of supply and resulting CGST/SGST vs IGST split will be confirmed with a tax adviser and stated on the invoice.** This is not asserted here as a settled tax position.

4.4 HSN / SAC codes — UNVERIFIED

The SAC codes below are marked UNVERIFIED and must be confirmed with a tax adviser before invoicing. They are not asserted as correct in this document.

SERVICE	SAC (UNVERIFIED — CONFIRM WITH TAX ADVISER)
Software development service	998314 (UNVERIFIED) — cited in CLIENT-ESTIMATE.md §7c; re-confirm t
AI / model training service	SAC: _____ (UNVERIFIED — BLANK. The correct SAC for AI/i

CLIENT-ESTIMATE.md §7c cites SAC 998314 for software development services at 18% GST. This annexure records that citation as **UNVERIFIED** for contractual use: the code must be confirmed against the current GST rate notification by the parties' tax adviser before it is printed on an invoice. For the AI/model-training services (the three fine-tuning workstreams), **no SAC is asserted at all** — the field is blank on purpose.

4.5 TDS

Under Indian law, this is a **domestic services contract**, and the client may be required to **deduct tax at source (TDS)** on payments to Vector Logic under the applicable section and rate for services. **The applicable section and rate must be confirmed with a tax adviser** and are not asserted here. Vector Logic's practical requirement: the **TDS deducted must be reflected on the client's TDS certificate (Form 16A) and credited to Vector Logic's PAN**, and the invoice will state the PAN so the credit can be matched. Where TDS applies, the client deducts it and pays the net; Vector Logic claims the credit. The contract should fix the TDS section, rate and certificate timeline explicitly.

4.6 Invoice mechanics Vector Logic requires

- A valid **tax invoice** with GSTIN, place of supply, SAC, invoice number and date.
- **GST charged on the invoice**, as the client's input credit where applicable.
- The invoice is raised on milestone delivery and acceptance, not on a calendar.
- PAN stated on the invoice for TDS matching.

5. SCOPE-CHANGE MECHANISM

5.1 How a change is raised and priced

Any change to the defined scope is raised in writing, described as a discrete deliverable, and priced before any work on it begins. Pricing is arithmetic on the engineering rate, shown to the client:

Engineering rate: \$35/hour = ₹3,385.84/hour = ₹27,087/day = ₹1,35,434 per 40-hour work-week (at the pinned FX rate of ₹96.7384/USD — see CLIENT-ESTIMATE.md §6).

A change order therefore quotes $\text{hours} \times ₹3,385.84$, or $\text{work-weeks} \times ₹1,35,434$, plus any pass-through vendor cost (GPU hours, studio time) at the then-current live price, plus GST.

5.2 The absorption threshold

Changes whose total value is at or below ₹13,542 (i.e. up to 4 engineering hours at ₹3,385.84/hour) are absorbed by Vector Logic at no charge. Changes above that threshold are **quoted as a change order**, approved in writing, and invoiced as delivered.

Why ₹13,542, i.e. 4 hours, is the defensible threshold: it is **half a work-day** — large enough to cover the small, real frictions that inevitably appear in a live build (a field added, a wording changed, a small integration tweak, a review call that turns into a fix) without a commercial conversation, and small enough that a genuine change of direction cannot hide inside it. Below half a day, disputing the change costs both parties more than the change; above half a day, the work is real enough that it must be quoted before it is done. **The threshold is a rate, not a yearly allowance** — it is applied per change, and a change cannot be split into sub-threshold pieces to avoid quoting.

5.3 What is not a change

Defect fixes within the warranty, corrections to delivered work that does not meet its own acceptance criteria, and work needed because Vector Logic mis-estimated its own effort are **not** change orders and are absorbed. Only a change to the agreed scope, or genuinely new scope, is quoted.

6. WHAT TRIGGERS THE MONTHLY CHARGE

The monthly charge of ₹1,26,108 begins only when the system is live and being used — that is, when it is serving real bookings on the client's behalf. It does not begin on handover date, and it does not begin while the system is idle waiting for the client's launch.

MONTHLY LINE	INR / MONTH
GPU serving — 1× 48 GB dedicated, 24/7	41,665
Managed Postgres + vector search	5,891
Redis	1,935
Real-time media server (SFU/TURN)	4,837
Object storage	1,161
Monitoring + CI	2,902
Support retainer (0.5 engineering work-week/month)	67,717
MONTHLY TOTAL	1,26,108

(Figures as in CLIENT - ESTIMATE .md §1b / the proposal's §6.2. The support retainer, ₹67,717, is deliberately larger than the GPU bill; a live 24/7 voice system needs ongoing engineering, not only hardware.)

6.1 The gap between handover and the first paying customer

This is a real commercial question and it has a plain answer:

- **In the gap, the system is not "live and being used" and the monthly charge does not run.** The

support retainer (₹67,717) does not accrue on an unused system.

- **Who pays the GPU in that window: the client, at cost, for the compute it chooses to keep warm.** The

GPU is a live vendor bill that falls due whether or not customers arrive. There are two honest options, and the client picks:

1. **Pause the serving instance** in the gap — a rented GPU can be stopped, and the client pays

nothing for compute while it is stopped. The system is brought up on launch. This is the cheapest and the recommended option where the gap is expected to be more than a few weeks.

1. **Keep a GPU warm** through the gap for pre-launch testing, demos and load rehearsal — in which

case the client pays the **GPU line (₹41,665/month) at cost**, with no support retainer and no margin for Vector Logic, for as long as it chooses to keep it warm.

- **Infrastructure that is already running** (managed Postgres, Redis, object storage, monitoring —

₹11,889/month together, being the non-GPU, non-retainer lines: $5,891 + 1,935 + 4,837 + 1,161 + 2,902 = 16,726$ — see the note below) is either paused or paid at cost, on the same principle as the GPU: the client pays for what it runs and nothing for what it does not.

Note on the arithmetic: the non-GPU, non-retainer monthly lines sum to **₹16,726** ($5,891 + 1,935 + 4,837 + 1,161 + 2,902$), and $₹41,665 + ₹16,726 + ₹67,717 = ₹1,26,108$, confirming the monthly total. In the gap, the only lines a client would pay at cost are the ones it leaves running; the retainer does not run.

In one sentence: the ₹1,26,108/month starts when the system serves real users; in the gap the client pays only for the infrastructure it chooses to keep running, at cost, and the retainer does not accrue.

7. TERMINATION

Stated fairly, in both directions.

7.1 If the client stops

- The client **pays for all milestones delivered and accepted** up to the date of termination, plus

any work in progress on a milestone it has approved, valued at actual engineering hours at ₹3,385.84/hour (or, where a milestone is substantially complete, at the milestone value on acceptance).

- **The mobilisation advance is non-refundable** where it has been spent on the front-loaded costs it

funded (GPU procurement and studio booking); where it has not been spent, the unspent portion is returned.

- The client **takes delivery of all work product paid for** — code, adapters, and the voice master, to

the extent they exist and are paid for.

- Any **vendor commitments already made on the client's behalf and non-cancellable** (a booked studio

session, a committed GPU period) are payable at cost.

7.2 If Vector Logic cannot deliver

- Vector Logic **refunds any advance received for scope not yet performed**, and delivers all work

product produced to that point.

- Vector Logic **transfers the trained model artefacts and the voice master** to the client, to the

extent they exist and the client has paid for them — the trained models are the client's property on handover, and a partial delivery transfers the partial artefacts the client has paid for.

- The **retention** (₹3,93,616) is not payable where handover acceptance has not been achieved, and any

retention already withheld is applied against sums owed.

- Where Vector Logic cannot deliver because a **Phase 0 gate failed**, this is **not** a Vector Logic

failure to deliver — it is the benchmark doing its job — and §3.9 governs, not this clause.

7.3 Trained model artefacts on termination

On any termination, the **trained adapters, the voice master, the training data ledger and the model evaluation results** are handed to the client to the extent they exist and have been paid for, **except that the voice master and any model trained**

on it are released only if the voice-rights release (§8) is in place. If no release exists, the voice-derived artefacts are withheld (see §8) until it does.

7.4 Termination for convenience

Either party may terminate for convenience on written notice. The financial consequences are the same as §7.1 where the client terminates, and §7.2 where Vector Logic terminates.

8. THE VOICE-RIGHTS RELEASE REQUIREMENT

The TTS brand voice is **recorded by Vector Logic in a studio session**, and it is a **priced line item** in this project — casting, session direction, QC, retakes and mastering (2 work-weeks, ₹2,70,868) plus the **studio recording pass-through (~₹97,000 for ~20 h of delivered audio) = ₹3,67,868**.

Because that recording is used to train a model that **becomes the client's property at handover**, an obligation follows, and it is not optional:

Before the recorded voice is used to train any shipped model, the client must provide a signed release from the voice artist granting perpetual, irrevocable, worldwide, transferable commercial rights to use that voice. The release must grant:

- **Perpetual** — no expiry;
- **Irrevocable** — the artist cannot withdraw it later;
- **Worldwide** — no territorial limit;
- **Transferable** — the rights can move with the business;
- **Commercial** — for commercial use, including training and serving a synthetic voice.

The release must name the client (Papparti) as the beneficiary, because the model is handed to the client and becomes the client's property; a release in anyone else's favour would leave the client's product exposed. The reason is concrete and commercial: **if the release were revocable, the client's product would break the day the artist withdrew** — a live 24/7 voice-first booking business cannot have its brand voice switched off by a third party.

Legal review of the release is a Phase 0 item (CLIENT-ESTIMATE.md §4 / DECISIONS.md D10).

8.1 What Vector Logic does if no release is provided

No release = no clone. If the client does not provide the signed release, Vector Logic **will not train the synthesis model on the recorded voice, and will not ship any model derived from it.** The recording, and any intermediate model trained on it, are held and not delivered. The affected workstream is **suspended, not proceeded with on the assumption that a release will arrive later**, and the client is notified in writing.

Financial consequence: where the synthesis workstream is suspended for want of a release, that portion of the model-work milestone is **not invoiced until the release is provided** (or, if the client elects not to proceed, it is removed from scope under §5 and credited). Vector Logic is not exposed to shipping a voice without rights, and the client is not billed for a deliverable that cannot lawfully be delivered. If the client elects to abandon the brand-voice workstream, the **voice-production and studio-recording costs actually incurred cancelling the booking remain payable**.

9. LIABILITY AND WARRANTIES

9.1 The accuracy limitations disclosed in the proposal survive into the contract

The accuracy limitations stated in §8 of the proposal survive into this contract and **no warranty in this contract is to be read as contradicting them.**

Specifically, and without limitation:

- **Hinglish code-switched speech recognition is not solved.** Residual word error is expected, and

no accuracy figure is warranted before Phase 0B measures one. Any figure in the proposal or the estimate is a research claim or a model-card number until Phase 0B produces a measurement.

- ****If the client supplies no real audio and no real transcripts, the real-world accuracy of the shipped**

models cannot be stated** — and therefore cannot be warranted.

- **Catastrophic forgetting is bounded and measured, never eliminated.** The contract warrants a

bounded and measured regression (target under 3%, measured before and after), not an absence of regression.

- **No listening test has been run and no MOS exists for the voice output** until the listening test is

run in the project. The contract does not warrant a quality score that has not been measured.

- **The avatar is a layered vector rig, not a photoreal likeness** — the proposal's §8.3 statement

is warranted as fact: what ships is what was shown, in kind.

9.2 What Vector Logic does warrant

- **It will measure before it asserts.** Phase 0B is run and reported against the gates in §3.1, and a

measurement is provided, not an opinion.

- **It will tell the client when a gate fails,** in writing, at Phase 0 cost — before the big spend.
- **It will deliver the artefacts it has accepted in the milestone schedule,** and they will be the

client's property on handover.

- **It will hold the source of truth for the numbers:** every figure traceable to the estimate, and the

engineering rate shown as arithmetic so it can be re-run.

9.3 What Vector Logic does not warrant

- **No accuracy figure is warranted before Phase 0B measures one.** In particular, no specific WER,

MOS, or tool-call accuracy figure is warranted at signature.

- **No vendor price is warranted to stay fixed** — for the reasons in §1, and consistent with the

validity window.

- **No business outcome** — bookings, revenue, take-rate profitability or market adoption — is warranted.

Vector Logic · Papparti

The proposal's §8.3 commercial risks (self-hosting is not cheaper at low volume; the commission take-rate is under threat; "any service" is the sector's failure mode) are the client's to weigh.

9.4 Bounded liability

Vector Logic's aggregate liability under this contract is capped at the total fees paid to Vector Logic under it. Liability for indirect, consequential, or loss-of-profit claims is excluded. This bounded, honest position is deliberate: an AI quality guarantee beyond measurement, or an uncapped liability on a project whose core accuracy question is not yet measured, would be a false comfort to both parties. What the contract offers instead of an open-ended warranty is **the Phase 0 gate mechanism** — measurable, cheap, and exercised before the big spend.

10. OPEN FIELDS REGISTER

Every field below is **BLANK ON PURPOSE**. Each is a fact that only a human — the client, Vector Logic's own records, or a licensed adviser — can supply. **None of them has been invented, and none may be filled by guesswork.** This annexure is not signature-ready until each is completed.

#	FIELD	WHO FILLS IT	MUST NOT BE INVENTED BECAUSE
1	Client legal entity name (registered name of Papparti's contracting entity)	Client	The correct legal person to bind is a client fact; no name has been supplied.
2	Client contact person + designation	Client	The authorised signatory is a client fact.
3	Client registered address (for the tax invoice)	Client	Required on a tax invoice; not supplied.
4	Client GSTIN (if registered)	Client	A registration number must be read from the certificate, never constructed.
5	Vector Logic GSTIN	Vector Logic	The registered entity's number must be read from its GST certificate.
6	SAC for software development service	Tax adviser (both parties)	The cited SAC is recorded as UNVERIFIED; the correct current code must be confirmed.
7	SAC for AI / model training services	Tax adviser (both parties)	No code is asserted. The field is blank. Confirm the correct SAC before invoicing.
8	Place of supply and CGST/SGST vs IGST split	Tax adviser	Determined under the IGST Act per invoice; not asserted here.
9	TDS section and rate	Tax adviser (both parties)	The applicable section and rate must be confirmed under Indian law.
10	Payment-days convention (e.g. Net 15 / Net 30)	Both parties	A commercial convention the parties agree; not assumed.
11	Late-payment interest rate	Both parties	Agreed in the contract; not assumed.
12	Acceptance-review window (days for deemed acceptance, §3.8)	Both parties	Agreed in the contract; not assumed.
13	Governing-law city / jurisdiction	Both parties	The seat of the contract is a negotiated choice.
14	Voice artist identity + release date	Client (with counsel)	A release is signed by a named artist; the identity and date are not known.
15	Effective date of the contract	Both parties	The date of execution.
16	Client PAN (for TDS matching, §4.5)	Client	A number to be read from the client's records.
17	Vector Logic PAN	Vector Logic	To be read from Vector Logic's records.
18	30-day warranty start date	Both parties	Fixed by the handover acceptance certificate (§3.7).

11. VERIFICATION

Every figure in this annexure was recomputed against `CLIENT-ESTIMATE.md` and the proposal before issue. The results:

1. **Payment schedule sum.** $7,87,232 + 4,47,967 + 23,03,239 + 16,25,205 + 10,83,470 + 10,83,470 +$

$5,41,735 = 78,72,318$ — exactly the one-time fee excluding GST. ✓

1. **GST arithmetic.** $₹78,72,318 \times 0.18 = ₹14,17,017$; $₹78,72,318 + ₹14,17,017 = ₹92,89,335$. ✓

2. **Phase-row consistency.** Phase 0 (4,47,967) + model work (30,90,471) + Phase 1 (16,25,205) +

Phase 2 (10,83,470) + Phase 3 (10,83,470) + deployment/handover (5,41,735) = **78,72,318**. ✓

1. **Monthly total.** $41,665 + 5,891 + 1,935 + 4,837 + 1,161 + 2,902 + 67,717 = 1,26,108$. ✓

2. **Engineering rate.** $\$35/\text{hr} \times 96.7384 = ₹3,385.84/\text{hr}$; $\times 40 = ₹1,35,434/\text{work-week}$. ✓

3. **Forbidden-token scan.** A case-insensitive scan of this document for the 17 forbidden base-model

tokens returned **0 occurrences of each** (see the tooling report accompanying this file).

What could not be verified and is marked accordingly:

- **SAC codes** — recorded UNVERIFIED; the software-development SAC is a citation pending confirmation,

and the AI/model-training SAC is blank.

- **Client identity, GSTIN, PAN, address, contact** — blank fields, §10.
- **Governing-law city, payment days, acceptance window, interest rate** — negotiated blanks, §10.
- **The studio recording cost (₹97,000)** is a modelled midpoint of a ₹64,000–₹1,30,000 band, not a

studio quote (`CLIENT-ESTIMATE.md` §10); it is stated as such where used.

Prepared by Vector Logic, 11 October 2026. This is a commercial document and not legal advice. All figures traceable to `CLIENT-ESTIMATE.md` and the proposal. Base model identities are commercially confidential to Vector Logic and are not disclosed.