

Dhwani

An AI-narrated audiobook publisher for vernacular Indian authors — Hindi, Tamil, Bengali, Marathi, Malayalam, Telugu, Gujarati — distributing to Audible, Storytel, Spotify, with the author's voice if they want it.

Category	Set 2 · India Services
Customer	Self-published and small-press Indian authors in regional languages; regional-language publishers wanting to expand into audio
Monetisation	60/40 royalty split (author/Dhwani) + ₹15,000–60,000 upfront production fee per book (waivable for high-WTP titles)
Build effort	Med
Plan version	v1.0 — 2026-05

Executive Summary

Dhwani is an AI-narrated audiobook publishing platform for the rapidly emerging Indian regional-language reading audience. The Indian audiobook market has grown ~45% per year since 2022, driven primarily by Audible's India expansion and Storytel's regional-language push; the great majority of this growth has been in Hindi and English content, leaving Tamil, Bengali, Marathi, Malayalam, Telugu, Gujarati, Punjabi, and Kannada audiences severely under-served. Regional-language audiobook production has been bottlenecked by the cost of professional human narration (₹35,000-1.2 lakh per hour of finished audio) and the difficulty of distributing through the major platforms.

Dhwani solves both bottlenecks. AI narration in Indian regional languages has reached broadcast-quality level in 2025-26 for narrative prose (slightly weaker on poetry and dramatic dialogue); Dhwani produces a finished audiobook from a manuscript in 2-4 weeks at a fraction of human-narration cost. Distribution is handled through Dhwani's pre-established channels into Audible, Storytel, Spotify, and Google Play Books. The business model: 60/40 royalty split (author keeps 60%) plus a one-time production fee that ranges from ₹15,000 to ₹60,000 depending on book length and complexity (waivable for high-WTP titles where Dhwani takes equity in the production).

Year-1 target: 800 audiobooks produced and published, generating ₹2.1 crore in revenue (combination of production fees and royalty share) against ₹95 lakh in costs. The wedge against existing options is dual: cost (vs. ₹2-7 lakh per book for human narration) and distribution access (vs. self-publishing where authors struggle to get into the major audio platforms). The wedge against pure AI tools (ElevenLabs + DIY distribution) is the production quality (curated narrator voices fine-tuned for each language, professional mastering, sound design) and the distribution pipeline.

The Problem

An Indian regional-language author — a Marathi novelist, a Tamil essayist, a Bengali short-story writer — wanting to publish their book as an audiobook today faces severe friction. Professional human narration in their language costs ₹35,000-1.2 lakh per hour of finished audio; a typical 280-page novel produces ~8 hours of audio = ₹3-10 lakh in narration alone, before mastering, editing, and distribution. For an author whose entire book advances ₹15,000-50,000 (typical for vernacular publishing), the audio cost is uneconomic.

Distribution is similarly difficult. Audible India accepts audiobooks only through approved partner imprints, not from individual authors. Storytel is more open but requires substantial volume commitments. Google Play Books is open but produces low discovery. The result is that vernacular audiobook content is dominated by the major publishers' backlists (Penguin Hindi, Manjul Prakashan, Bharatiya Bhasha Parishad) and a small set of celebrity authors. The mid-tier and self-published vernacular author has no path.

The audience exists and is growing. Storytel India's regional-language listenership has grown ~70% year-on-year for 4 consecutive years. Tamil and Malayalam audiobook demand on Audible India is in the top quartile globally. The market wants the content; the supply side cannot economically produce it.

The Solution

Dhwani offers a productised audiobook publishing service. Author submits manuscript (DOC or PDF), selects narrator voice from a curated library of 25-50 voice options per language (each fine-tuned for narrative reading with appropriate tonal register), reviews the AI-produced sample chapter, then approves full production. Within 2-4 weeks, a mastered audiobook is produced and submitted to Audible, Storytel, Spotify, and Google Play Books through Dhwani's distribution pipeline.

Three structural differences from existing options define the wedge. First, cost: ₹15,000-60,000 production fee + 40% royalty share (vs. ₹3-10 lakh for human narration with the same royalty share). Second, distribution: Dhwani has established imprint relationships with Audible India and Storytel India that solo authors cannot replicate. Third, language depth: the narrator-voice library is curated for each of 8 languages with native-speaker tonal calibration, not generic TTS.

Author Voice option: for authors who want their own voice narrating (a meaningful share of memoir and non-fiction authors), Dhwani offers a voice-cloning workflow. The author records 90 minutes of clean reference audio in a guided session; Dhwani trains a voice model on the recording (~3 days); the AI then narrates the full book in the author's own voice. Premium production fee: ₹85,000.

Premium tier (₹85,000+ for top-WTP authors): full production with custom sound design (subtle music cues, ambient sound), human editor pass on the AI narration for nuance correction, and accelerated 7-day turnaround. Suitable for serial-publishing authors and small-press partners.

Market Opportunity

The Indian audiobook market reached approximately ₹950 crore in 2025 across all platforms and languages, growing at ~35-45% per year. Regional-language share of this is currently ~22% but growing faster (~55% year-on-year) than English (~25%). By 2028, regional-language audiobook market is projected at ₹1,200-1,800 crore.

The supply side: there are approximately 18,000 actively-published vernacular Indian-language authors with at least one book to their name across the 8 major languages. Of these, an estimated 6,000 have published in the last 3 years (active authors). Audiobook-conversion penetration is currently below 8% for this set. The addressable supply opportunity is therefore 5,500+ authors who could potentially produce audiobooks if economics worked.

At an average lifetime value of ₹85,000 per book (production fee + 40% royalty share over 5 years for a moderately successful audiobook), the supply-side SAM is approximately ₹470 crore. Demand-side growth makes this number compound favourably.

Adjacent opportunities: short-form audio (podcast-length content from established writers, monetised through Spotify and Storytel), backlist audio (older books from publishers wanting to monetise the long tail), and audio-only original content (commissioning new audiobooks for genres that work well in audio).

Target Customer

Primary persona: a 52-year-old Tamil novelist with 7 published novels, of which 4 are out of print. Sales of her current titles ~3,000-8,000 copies per book across 10 years of life. Would love to have her work in audio (her readers are aging; some have asked) but has been quoted ₹4 lakh by a professional narration studio. Will pay Dhwani's ₹35,000 production fee for her best-selling title without negotiation, with serious interest in expanding to her backlist if the first audiobook performs.

Secondary persona: a 41-year-old Bengali non-fiction author writing on Indian political history. Has 3 active books. Wants to narrate them in his own voice (his readers know his voice from interviews and book launches). Will pay ₹85,000 Premium Voice tier for the voice-cloned narration.

Tertiary persona: a 38-year-old small-press founder publishing 12-25 vernacular titles per year across Marathi, Hindi, and Gujarati. Wants audio for the full catalog but production costs make it impossible on per-book basis. Will enter into a volume agreement with Dhwani: 15 books per year at preferential ₹22,000 per-book production fee + 50/50 royalty split (better split for higher volume).

Product

Author intake: web-based onboarding (manuscript upload, language selection, narrator-voice selection from curated library, basic metadata). 25-minute video orientation call with a Dhwani producer for first-time authors.

Voice library: 25-50 curated narrator voices per language, each fine-tuned on professional narrator recordings (Dhwani has licensed voice data from professional narrators in each language, with revenue-share or buyout structures). Voices categorised by tonal register: formal, casual, intimate, dramatic, regional dialect variants.

Production workflow: manuscript parsed into chapters → AI narration produced with chosen voice → automated mastering (level adjustment, noise removal, chapter markers) → human QA review on a 10% sample per book → final mastered file. Standard turnaround 14-28 days.

Voice-cloning option (Premium): author records 90 minutes of guided reference audio (read from provided neutral text, in a quiet setting with provided microphone if needed). Voice model trained on recording (~3 days). Full book narrated in cloned voice with same downstream production.

Distribution: pre-established pipelines into Audible India (via Dhwani's imprint relationship), Storytel, Spotify Audiobooks, Google Play Books, and Apple Books. Author receives consolidated monthly royalty report across all channels.

Author dashboard: production status, performance analytics (downloads, listens, royalty earned per channel), reader reviews aggregated across channels.

Technical Architecture

Frontend: Next.js 14 + Tailwind for author portal.

Backend: Python FastAPI on Hetzner cloud. Postgres on Neon. S3-compatible storage for audio files (substantial — typical audiobook is 200-500 MB).

AI narration stack: hybrid. For Hindi/English: ElevenLabs Multilingual + custom voice models (~₹0.60 per minute of audio output). For regional languages: AI4Bharat IndicTTS as base + custom voice fine-tuning per language on professional-narrator reference data. Open-source TTS infrastructure runs on Hetzner GPU box (₹35k/month for sustained capacity).

Voice cloning: open-source XTTS v2 + custom training pipeline. ~₹3,500 in GPU cost per voice trained.

Audio processing: FFmpeg + Sox for mastering, automated chapter-marker insertion, level normalisation.

Distribution integrations: APIs/SFTP feeds into Audible (via ACX submission), Storytel (direct relationship), Spotify Audiobooks (via Findaway Voices intermediary), Google Play Books (direct), Apple Books (via Author Direct).

Business Model & Unit Economics

Two revenue streams. Production fee (upfront, per book): ₹15,000-60,000 standard, ₹85,000 Premium with Voice cloning, with volume discounts for publisher partners (~30% reduction at 15+ books/year commitment). Royalty share (ongoing): 40% of net royalties from all distribution channels accrues to Dhwani, 60% to author (the standard 60/40 split skewed toward author).

Conversion economics: sales cycle 4-8 weeks (author decision; for publishers 6-12 weeks). Conversion rate from qualified inbound author: 35%. Distribution: 65% standard production, 25% Voice cloning Premium, 10% publisher-volume tier.

Per-book gross margin: Standard production at ₹35,000 fee — direct AI/audio cost ~₹2,800, QA labour ~₹4,500, distribution overhead ~₹1,200 = ~₹8,500 direct cost = 76% gross margin on production fee. Royalty share is ~92% gross margin (mostly platform pass-through with minimal direct cost). Voice cloning Premium has lower margin on the per-book fee (additional GPU cost + extended QA) but typically generates higher royalty share due to higher author engagement and promotion.

Unit Economics (Year-1 base case)

Year-1 audiobooks produced (target)	800
Year-1 revenue (production fees + Year-1 royalty share)	₹1.2 crore
Gross margin	78% blended
Author acquisition cost	₹4,200
Payback period	Single production cycle
Year-1 all-in costs	~₹95 lakh
Year-1 net contribution	~₹1.0 crore
Royalty share tail (years 2-5 per book)	₹6,000-15,000 average

Go-to-Market

Channel 1 — Direct outreach to vernacular authors and small publishers (40%): Indian vernacular literary networks are densely connected through regional literary festivals (Jaipur Literature Festival, Chennai Book Fair, Kolkata Book Fair, regional sahitya parishads); direct outreach via these networks reaches authors directly.

Channel 2 — Publisher partnership programme (25%): partnerships with 20-40 small/mid-tier regional publishers for volume audio production at preferential pricing. Publishers gain audio offerings without operational investment; Dhwani gains volume.

Channel 3 — Successful author case studies + word of mouth (20%): the first 50 high-quality audiobooks become case studies that drive subsequent author inbound. Strong word of mouth in tightly-networked vernacular literary communities.

Channel 4 — Content + literary-event presence (15%): sponsorship and presence at major regional literary events, plus content (a blog covering vernacular audiobook trends, narrator-of-the-month profiles, etc.) for ongoing author engagement.

Roadmap (first 12 months)

- Month 1-3: Voice library curation for Hindi + Tamil + Bengali + Marathi (the four highest-demand languages), distribution pipeline setup with Audible + Storytel, production workflow MVP, first 30

audiobooks produced.

- Month 4-5: Expand to Malayalam + Telugu + Gujarati + Punjabi + Kannada, 130 audiobooks produced cumulative, ■35 lakh annualised revenue.
- Month 6-8: Voice cloning Premium tier launched, first publisher-partnership deals (2-3 publishers), 350 audiobooks cumulative.
- Month 9-10: Spotify and Google Play distribution pipelines added, expanded narrator library, 550 audiobooks cumulative.
- Month 11-12: 800 audiobooks cumulative, ■2.1 crore revenue, 6 active publisher partnerships.

Key Risks

- Voice quality plateau or backsliding: AI TTS in regional languages has improved rapidly but is not yet flawless for poetry or dramatic dialogue — mitigated by clear authoring expectations (we excel at narrative prose, weaker at heavily-dialogue or poetic content), by sample-chapter approval gates before full production, and by human-editor pass option for higher-WTP titles.
- Audible India policy changes: Amazon's ACX programme could change submission rules or compensation structure — mitigated by multi-platform distribution (no single platform >50% of revenue), by Storytel as a counterweight, and by direct-author-website distribution as a fallback for difficult cases.
- Author resistance to AI narration on principle — meaningful share of literary authors view AI narration as artistically inferior or threatening to human narrators — mitigated by transparent positioning (we are the option for authors who cannot afford human narration, not a replacement for those who can), by Voice cloning option (author's own voice), and by the human-editor-pass tier for authors who want it.
- Royalty-share collection complexity across 5 platforms each with quarterly settlement and varying rules — operational burden; mitigated by automated royalty-reconciliation tools and by a dedicated finance person from Month 6.
- Reputational risk from a single low-quality AI-produced audiobook going viral as 'this is bad' — mitigated by 10% sample QA on every production, by author-approval gate before publication, and by clearly opt-out language permitting an author to reject the final mastered file without penalty.