

# VerticalScribe

*AI scribes for the underserved professionals — immigration lawyers, veterinarians, property surveyors, court reporters — where medical scribes are saturated and these aren't.*

<b>Category</b>	Set 3 · Post-AI Plays
<b>Customer</b>	Single-vertical professionals doing structured documentation: immigration lawyers, veterinarians, court reporters, property surveyors, social workers, structural engineers
<b>Monetisation</b>	\$199–449/mo per seat (varies by vertical) · \$5k–25k/mo enterprise multi-seat
<b>Build effort</b>	Med
<b>Plan version</b>	v1.0 — 2026-05

## Executive Summary

VerticalScribe is a focused AI-scribe product targeting the professionals beyond medical practice — where the medical-scribe market has become saturated and competitive (Abridge, Nuance DAX, Suki, Heidi, Freed all chasing the same 800k US physicians) and where adjacent professional categories are equally documentation-heavy but completely unserved. The opportunity: immigration lawyers writing detailed asylum-application narratives, veterinarians producing post-visit summaries, property surveyors generating structural assessment reports, court reporters delivering certified transcripts, social workers documenting case notes, structural engineers writing inspection reports — each profession produces 30-90 minutes of written documentation per client interaction, all amenable to AI-scribe augmentation.

The product takes the same form as medical scribes — record the meeting or dictation, AI generates a structured first draft in the profession's required format, human reviews and edits — but tuned specifically for each vertical's compliance, format, and tone requirements. Pricing is \$199-449/month per seat depending on vertical (immigration lawyers \$349-449 given high billable-hour value; veterinarians \$199-249; court reporters \$399-449), with enterprise multi-seat pricing at \$5k-25k/month for firms.

Year-1 target: 1,400 paying seats across 4 launched verticals generating \$4.2 million revenue (~₹35 crore) against ₹11 crore in costs. The wedge against horizontal AI tools (ChatGPT, Claude) is the vertical-specific output format and compliance language baked in (an immigration lawyer's narrative must follow specific USCIS conventions; a veterinarian's SOAP note must match veterinary documentation standards; a property surveyor's report must address state-specific structural-assessment requirements). The wedge against medical-scribe incumbents is going one inch wide and a mile deep into underserved verticals they will not prioritise.

## The Problem

Documentation is the single largest non-billable time cost for most professional service practitioners — equivalent to 30-45% of working hours for the typical immigration lawyer, veterinarian, court reporter, surveyor, or social worker. Medical practice has solved this through dedicated AI-scribe products in 2023-25 (Abridge, Suki, Nuance DAX, Freed); the bulk of US physicians now use one. But the technology has not migrated to adjacent professions despite identical underlying need.

Each underserved profession has its own documentation pattern. An immigration lawyer preparing an asylum application writes a 12-25 page declaration capturing the client's persecution narrative, with specific structural requirements (chronological, named-actor, evidence-linked). A veterinarian documents each appointment in SOAP format (Subjective, Objective, Assessment, Plan) integrating clinical observations, owner-reported history, exam findings, and treatment plan. A property surveyor produces a 8-30 page structural assessment with state-specific format requirements covering foundation, framing, roof, systems. A court reporter transcribes proceedings to specific certification standards.

The horizontal AI tools (ChatGPT, Claude, Microsoft 365 Copilot) can produce a generic first draft from a recording but miss the profession-specific formatting, compliance language, terminology, and structural conventions that distinguish a usable professional document from a generic AI output. The professional ends up spending nearly as much time editing the AI draft into proper form as they would have spent writing it from scratch. The vertical-specific scribe — which incorporates the format requirements, terminology, and conventions natively — produces 70-85% time savings, the same magnitude as medical AI scribes have demonstrated.

## The Solution

VerticalScribe is a focused product offering AI scribing for one vertical at a time. Each vertical launch involves: deep partnership with 8-15 expert practitioners to learn documentation conventions, format requirements, compliance language; fine-tuning of a vertical-specific output layer on top of base LLMs; clinical / professional validation; pricing calibration; vertical-specific go-to-market.

Product flow per vertical is consistent. Practitioner records the client interaction (in-person consultation, phone, video call, or dictation) via mobile app, web app, or browser extension. Recording uploaded; AI produces a structured first draft in the profession's specific format within 60-180 seconds depending on length. Practitioner reviews, edits in clean web editor with vertical-appropriate templates and citations, finalises. Document exported to the practitioner's workflow tool (case-management system for lawyers, practice-management software for veterinarians, etc.).

Three structural differences from horizontal AI tools and from medical-scribe incumbents define the wedge. First, vertical-specific format: output is in the exact structural form the profession requires, with profession-specific terminology and compliance language. Second, integration depth: integrates with the case-management or practice-management software each vertical uses (Clio for lawyers, AVImark/Cornerstone for veterinarians, etc.) rather than producing standalone output. Third, focused investment in underserved verticals that medical-scribe incumbents will not prioritise.

Launch sequencing matters. Vertical 1: immigration lawyers (highest seat ARPU, clear ROI math, accessible community via AILA). Vertical 2: veterinarians (large market, clear documentation pattern, accessible via VMA chapters). Vertical 3: court reporters (high ARPU, regulatory tailwind as remote court grows). Vertical 4: property surveyors (high ARPU, fragmented market that no horizontal tool serves well). Year 2 adds 3-4 more verticals.

## Market Opportunity

Each addressable vertical has a finite but meaningful target seat count. Immigration lawyers (US): 15,000 active practitioners. Veterinarians (US + Canada + Australia + UK): 240,000 practitioners. Court reporters (US): 28,000 active. Property surveyors (US + Canada + UK): 110,000 active. Social workers in private practice: 180,000 (US). Structural engineers in private practice: 45,000. Combined launch-tier addressable market: ~620,000 seats.

At a blended ARPU of \$3,600/year per seat (across the verticals), the SAM is approximately \$2.2 billion. Realistic 5-year capture target: 1-2% of total addressable seats ≈ 6,000-12,000 paying seats = \$22-44 million ARR. Comparable trajectory to medical-scribe companies in their first 3-4 years (Abridge reached \$50M ARR in year 4).

Adjacent verticals for year 3-5 expansion: dentists (separate from medical scribes in workflow), tax accountants, real-estate appraisers, insurance claims adjusters, mental-health therapists in private practice. Each is a \$40-200M sub-market with similar dynamics.

## Target Customer

Primary persona: a 44-year-old immigration lawyer in Houston with a 4-attorney firm. Personally drafts 25-40 asylum and visa application declarations per month; each requires 90 minutes to 4 hours of writing. Currently spends 50-60% of working hours on documentation; bills only the other 40-50%. Will pay \$449/month for VerticalScribe immigration-lawyer tier; recovered billable hours pay for the subscription many times over.

Secondary persona: a 36-year-old veterinarian at a 6-doctor practice in Denver. Sees 22-28 patients per day; each requires a SOAP note typically written between patients or after close. Documentation backlog at end of day is the practice's biggest source of staff frustration and lost billing. Will pay \$229/month for VerticalScribe veterinary tier; the multi-doctor practice will adopt for all 6 doctors at firm-tier pricing (\$1,099/month total).

Tertiary persona: a 52-year-old court reporter in California producing certified transcripts for depositions and court proceedings. Spends 2-3 hours of post-proceeding documentation per hour of proceeding. Will pay \$399/month for VerticalScribe court-reporter tier with the certified-transcript output template.

## Product

Multi-modal capture: native mobile app (iOS + Android) for in-person recording, web app with built-in recording for desk-based sessions, browser extension for video calls (Zoom, Google Meet, Microsoft Teams), dictation mode for solo-dictated notes.

Vertical-specific output layer: each vertical has a dedicated fine-tuned prompt structure, format template, terminology dictionary, compliance language library, and citation/reference patterns. Output is in the exact form the profession requires.

In-product editor: clean web editor with vertical-appropriate templates, side-by-side recording-and-text view, one-click insertion of standard clauses/sections, version history, comment threads for team review.

Integration layer (per vertical): immigration lawyers integrate with Clio + LollyLaw; veterinarians integrate with AVImark, Cornerstone, Provet, ezyVet; court reporters integrate with CaseCAT, eclipse, court-management systems; surveyors integrate with InspectIt, HomeGauge, Spectora.

Team / firm features: multi-seat workspace, shared template libraries, supervisor review queues (especially valuable for immigration firms where senior attorneys review junior work), client folders, billing-time tracking integration.

## Technical Architecture

Frontend: Next.js 14 web app + React Native mobile + browser extensions. Real-time recording UI with waveform display and live transcription preview.

Backend: Python FastAPI on Hetzner cloud, scales to thousands of seats before sharding required.

Transcription: hybrid approach. Faster-Whisper on Hetzner GPU boxes for English (cost ~\$0.012/minute, vs. OpenAI Whisper API at ~\$0.06/minute). OpenAI Whisper API as fallback for non-English languages and complex audio.

Vertical-specific generation: GPT-4o for English output with vertical-specific prompts and templates (~\$0.18 per typical document). Anthropic Claude as secondary provider for tone-sensitive outputs.

Integration platform: per-vertical integration code with major practice-management / case-management systems. ~3-6 weeks engineering work per integration; integration moat compounds as we ship.

Compliance per vertical: medical-style HIPAA-aligned controls for veterinary (similar requirements), immigration-specific privilege protection for legal, audit logging for court reporter use cases.

Payments: Stripe (USD primary), with multi-currency support for international expansion.

## Business Model & Unit Economics

Per-seat per-month pricing, with per-vertical price points. Immigration lawyer: \$349-449/seat/month (high ARPU justified by ROI math). Veterinarian: \$199-249/seat/month. Court reporter: \$399-449/seat/month. Property surveyor: \$279-329/seat/month. Firm/practice tiers (5+ seats): 25% per-seat discount. Enterprise (50+ seats): negotiated \$5k-25k/month.

Conversion economics: free trial (14 days, fully functional) converts at 22%. Annual prepay (15% discount) adopted by 40% of seats; reduces churn meaningfully. Monthly churn target under 3.5% on monthly; under 1.5% on annual.

Gross margin: ~78% blended. Major cost lines: transcription (~\$11/seat/month at average usage), LLM generation (~\$8/seat/month), per-seat infrastructure (~\$3/seat/month).

Customer LTV per seat: \$3,600 × 2.5-year average lifetime = \$9,000 LTV. CAC target: \$850. LTV/CAC: 10.6.

### Unit Economics (Year-1 base case)

<b>Year-1 paying seats (target)</b>	1,400 (across 4 verticals)
<b>Blended ARPU per seat</b>	\$3,000/year
<b>Year-1 revenue</b>	\$4.2 million (~₹35 crore)
<b>Gross margin</b>	78%
<b>Customer acquisition cost (CAC)</b>	\$850
<b>Payback period</b>	3.4 months
<b>Year-1 all-in costs</b>	~₹11 crore
<b>Year-1 net contribution</b>	~₹16 crore

## Go-to-Market

Channel 1 — Vertical-association partnerships (40%): partnerships with AILA (immigration), AVMA (veterinary), NCRA (court reporters), NSPS (surveyors) for member discounts, sponsored webinars, conference presence. Highest-quality channel for professional services.

Channel 2 — Targeted content for each vertical (30%): substantive content addressing each vertical's specific documentation pain points (asylum-narrative best practices, SOAP-note efficiency, certified-transcript turnaround). Build inbound from in-vertical SEO and from association-newsletter syndication.

Channel 3 — Practice-management vendor partnerships (20%): partnerships with Clio (legal), AVImark / Cornerstone (veterinary), CaseCAT (court reporting) for integration showcases and joint go-to-market.

Channel 4 — Practitioner-referral programmes (10%): per-vertical referral incentives. Professional service practitioners refer well within their tight communities.

### Roadmap (first 12 months)

- Month 1-3: Vertical 1 (immigration lawyers) — partnership with 12 expert immigration attorneys, vertical-specific prompts and templates, integration with Clio + LollyLaw, launch with first 80 seats.
- Month 4-5: Vertical 1 scale to 220 seats, AILA partnership signed, Vertical 2 (veterinary) development begin.

- Month 6-8: Vertical 2 (veterinary) launch with AVImark / Cornerstone integrations, Vertical 1 at 380 seats, Vertical 2 at 250 seats.
- Month 9-10: Vertical 3 (court reporters) launch, Vertical 1 at 480, Vertical 2 at 450, Vertical 3 at 90.
- Month 11-12: Vertical 4 (property surveyors) launch, total seats 1,400, \$4.2M ARR.

## Key Risks

- Medical-scribe incumbents (Abridge, Suki) expanding into adjacent verticals as their core market saturates — possible 24-36 month threat; mitigated by vertical-depth investment that incumbents will not prioritise (each adjacent vertical has documentation conventions requiring multi-month domain learning) and by first-mover association partnerships that create switching cost.
- Horizontal AI tools (ChatGPT, Claude) becoming good enough at vertical-specific output that the wedge disappears — partial threat; mitigated by integration depth with practice-management systems (horizontal tools cannot easily integrate), by vertical-specific compliance language requirements, and by the workflow advantage of a focused product over a chat interface.
- Vertical-specific compliance variations across jurisdictions: immigration law differs USCIS vs. UK Home Office; veterinary differs state-by-state; surveying has 50 different state codes — operational complexity scales with jurisdictional expansion; mitigated by jurisdiction-by-jurisdiction rollout and by reusable compliance-language infrastructure.
- Slow per-vertical onboarding: each vertical takes 2-3 months of expert partnership to launch credibly; this constrains expansion speed — mitigated by parallel vertical-launch tracks (different team members on different verticals) and by accepting the trade-off (deep launch beats shallow expansion).
- Professional liability if scribe-generated output causes harm (incorrect legal narrative, missed clinical observation) — substantial risk; mitigated by professional indemnity insurance, by clear scope-of-product disclaimers ('AI draft requires practitioner review and approval'), and by audit logging of every edit / sign-off.