

LansonAI

The Voice Context Layer for Live Speech

Starting with Lanson Live: stable real-time captions people can actually follow

Real-Time Captions

Stable captions as speech unfolds
Zero reflow. Zero re-anchoring.

Stable Real-Time Captions

Attention Anchor keeps text from
moving

Patent Family Moat

Cognition-safe live caption delivery,
deployed in production

Founder

Zhen Zhang

Ex-Tencent (5 years, 100M+ DAU systems)

Deal Terms

Raising: Pre-seed SAFE (post-money)

Initial allocation: **\$100K @ \$10M** cap

Remaining allocation: **\$200K @ \$12M** cap

Target: **\$300K** (hard cap **\$500K**)

Problem: Live Speech Is Hard to Consume

Live speech is hard to read, follow, reuse, and pass to AI.

When you type, you process as you go: structure emerges naturally as you choose words, revise, and refine in real time.

When you speak live, words are linear, fleeting, and unstructured — hard to read in real time, hard to follow, and hard to pass to AI.

It's not just that speech becomes messy after the fact. The deeper issue is that live speech has never been designed to be consumed as it happens.

Voice today still lacks a reliable live context layer: a caption stream that turns speech into something people can follow instantly and systems can use immediately.

Speech is natural to produce, but still hard to consume as live context.

Why Now: The Inflection Point Is Here

Two forces converged simultaneously — **for the first time in history.**



LLMs Changed the Foundation

- For 30 years, voice failed not because machines couldn't hear — but because they couldn't understand context.
- LLMs reached production-grade reliability in 2023. The missing context layer now exists and is deployable at scale — enabling **context-aware refinement.**



Spatial Computing Is Arriving


- Apple Vision Pro and Meta Quest have shipped. Stable, cognition-safe captions are no longer a feature — they are baseline infrastructure.
- For the first time, a hardware platform makes live caption quality **mandatory infrastructure.**

The Window Is Open — Right Now

The window opened in 2023–2024. The race to define the live caption layer as foundational infrastructure is happening now — and early-mover position creates compounding defensibility technically, legally, and in distribution.

Trinity Engine: The Post-Inference Layer

Trinity Engine controls the post-inference layer that turns raw speech into stable, readable, real-time context.



Performance Metrics

120×*

Speed

1 hour audio → 30 seconds

90%*

Zero-Edit Rate

Publication-ready output

80%*

Cost Reduction

Significant cost savings

† 120× speed | 90% zero-edit rate | 80% cost reduction are based on internal benchmarks. Full methodology and test conditions in Appendix.

Stable Captions Keep the Text from Moving. Attention Anchor Keeps the Eyes from Searching.

Technology sets our ceiling — product philosophy determines how far we go.

How Live Captions Work Today



Flow-Layout Rendering

Text blocks reflow constantly, causing users to lose their place and experience motion sickness.



Cognitive Re-Anchoring

Constant text shifting forces the brain to re-anchor, creating a fundamental barrier to following live speech.

Stable Real-Time Captions Deliver



Attention Anchor

Caption containers use stable coordinates so text appears in place without reflowing. Users stay anchored to the speaker.



Cognition-Safe Delivery

Follow live conversations naturally without re-anchoring or motion sickness. **Try it yourself** — link below.

What our patent protects

Not a feature — the core method of delivering live speech context without cognitive disruption. Already deployed and measurable.

Lanson Live (no login, limited to 5 mins)

realtime.lansonai.com

Lanson Live (Google/Apple 1-click)

live.lansonai.com

Patent Strategy: Early-Mover Advantage

What This Patent Covers

- Establishes prior art and first-to-file position on cognition-safe live caption delivery (Zero-Reflow Rendering)
- Creates licensing negotiation leverage across 2D, AR, and spatial computing surfaces
- Forces competitors to design around our method — or license it
- Already deployed in production: this is not a theoretical claim
- The patent is not about showing captions. It protects a method for delivering live speech context without forcing users to visually re-anchor as captions update.

Why Platforms Need This

- Apple Vision Pro: Accessibility APIs becoming mandatory infrastructure
- Meta Quest: Real-time captions required for key enterprise and social use cases
- AR Glasses Wave: Stable, cognition-safe captions are table stakes for mass adoption
- Every major platform entering spatial computing will need a solution — we already have one

Licensing Potential

- Comparable IP deals in voice/display suggest meaningful per-device royalty potential
- Platform integration licenses represent a separate revenue layer from SaaS
- Strategic acquisition optionality: 3–5 year horizon
- Patent grant expected; licensing is upside, not base case

We welcome technical due diligence at any level of scrutiny — and we'd love to show you what's already running in production. If you're building in spatial computing, accessibility, or voice infrastructure, let's talk.

Market: Bottom-Up TAM with Clear Path

Lanson Live	Lanson Podcast	API & SDK	Platform / IP
Today / Primary	Today / Secondary	Year 1+	Long-term
Meetings, lectures, interviews, live events, accessibility, creators / livestreams	5M monetized creators (>\$1K/month)	Usage-based pricing for B2B integrations	AR/MR, accessibility infrastructure
Pricing: subscription (comparable to or higher than Podcast tier)	Pricing: \$49.99/mo (Creator Pro)	Real-time captions, real-time translation, caption stream / caption layer	Voice Context Layer for AI systems
Year 3 target: 0.5% of addressable live speech use cases	Year 3 target: 0.5% penetration ≈ \$14.7M ARR	Target: conferencing tools, media platforms, accessibility layers	Patent licensing to Apple, Meta, Google OEMs Upside optionality, not base case

Revenue Scenarios

Base Case	Bull Case	Platform Case	Lottery Ticket
Lanson Live + Podcast subscriptions	+API/SDK real-time caption infrastructure layer	Enterprise SAM +\$8B addressable	AR/MR IP licensing

Organic Paid Signal Before Distribution

Lanson Live generated organic paid conversion shortly after App Store launch — before formal launch, paid ads, or proper ASO optimization.

127

Active Users

100% organic beta

523

Hours Processed

Total platform usage

90%

Zero-Edit Rate

vs. 34–41% industry avg

76%


Willing to Pay

From user survey


Retention & Usage

- Wave 1 (Oct 2025): 34% retention
- Wave 2 (Dec 2025): 58% retention
- Heavy Users (10+ hrs): 8 users
- Average: 4.1 hrs/user | Median: 2.3 hrs/user

Quality Benchmark (Zero-Edit Rate)



Whisper raw



Otter.ai



Descript

01



LansonAI

02

Month 1 Launch Plan

Scale to 30-40 paying users at \$49.99/mo

Month 3 Target

\$5K MRR (~100 paying users at \$49.99/mo)

Competitive Landscape: A Different Category

Dictation is for input. Transcription is for records. Meeting AI is for summaries. Lanson Live is for following speech as it happens.

Category Comparison

Category	Representative	User Mental Model	Lanson Live Difference
Meeting memory	Otter.ai	Post-meeting records / summaries	Following speech as it happens
Editing tool	Descript	Post-production editing	Live captions, not editing
Raw STT API	Deepgram / AssemblyAI	Words / API output	Stable readable caption layer
Caption tools	Browser caption tools	Single-point captions	Caption stream + context-aware refinement + sharing + IP
Lanson Live	LansonAI	Live speech easy to follow	Stable real-time captions

Lanson Podcast vs. Content Tools

Company	Speed	Quality	Pricing
Otter.ai	✓ Real-time	34% zero-edit	\$16.99/mo
Descript	✓ Real-time	41% zero-edit	\$24-35/mo
LansonAI	Up to 120× faster*†	90% zero-edit*†	\$49.99/mo

† 120× speed and 90% zero-edit rate are based on internal benchmarks.

Full methodology and test conditions in Appendix.

Our Unfair Advantages

- **Trinity Engine™**: post-inference layer no competitor has
- **Patent family**: Cognition-safe live caption delivery (Zero-Reflow Rendering)
- **IP License**: first-mover position across 2D, spatial computing, and future HCI surfaces
- **Execution speed**: Shipped production system in 6 months — iOS / Android / Web
- **Infrastructure cost**: < \$0.0005/audio min — **~10× unit-economics headroom** vs. public market list pricing

Team: Technical Founder with Billion-Scale Experience

5 years at Tencent — building systems serving 100M+ DAU across QQ and WeChat ecosystems

What Makes This Founder Different

1

Systems Architecture at Scale

Five years at Tencent building and maintaining infrastructure serving 100M+ DAU

2

End-to-End Product Execution

Built Web, iOS, and Android — solo — in 6 months. Everything included, end to end.

3

Taste and Product Judgment

Built and refined the product without a design team, PM, or marketing budget.

What I've built



Full-stack product (Web / iOS / Android)



Trinity Engine processing pipeline



Serverless infrastructure with control center



Lanson Live: organic paid conversion before formal launch or paid ads



Patent design, written, filed and prosecuting




Brand identity, positioning & launch video

Solo founder risk

The strength and the risk come from the same place. I have been able to build with unusual coherence and speed because key decisions have stayed tightly integrated. After the round, I will strengthen the company with senior hires in growth and engineering.

Business Model: Product → Platform → IP

Primary: Lanson Live subscription



Unit economics (projected):

\$50

CAC

\$15

COGS

5%

Churn


\$700

LTV


* LTV calculated based on Focus tier (\$49.99/mo), \$15 COGS, and 5% monthly churn (industry benchmark). Blended ARPU across tiers TBD.

All paid plans include real-time translation. Higher tiers unlock broader language coverage across major global languages.

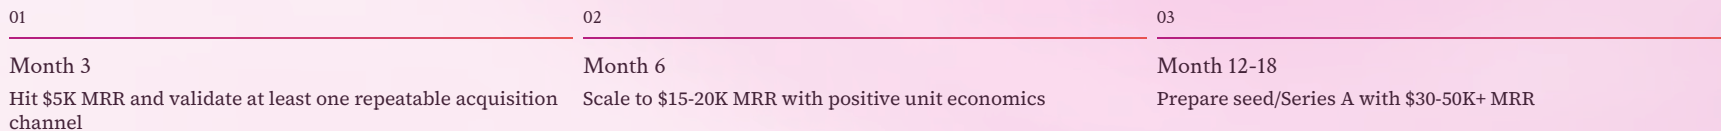
Secondary: Infrastructure — API / SDK



Tertiary: Patent licensing



12-18 month milestones



The Ask: Pre-Seed SAFE (Rolling Close)

The product, technical foundation, and market narrative are already in place. We are raising because growth is the next constraint.

Deal Terms at a Glance

- **Instrument:** Post-money SAFE
- **Initial Allocation:** \$100K @ \$10M cap (for early believers)
- **Remaining Allocation:** \$200K @ \$12M cap
- **Target Raise:** \$300K | **Hard Cap:** \$500K
- **Structure:** Rolling close (investors can sign and wire individually)

Use of funds (6-9 months)

Growth & Distribution (50%)

Distribution for Lanson Live's high-intent use cases: ASO / SEO, accessibility communities, live event demos, creator / livestream workflows, education / lecture use cases, onboarding and paywall experiments

Product & Engineering (30%)

Improving Trinity Engine™, caption stabilization, real-time translation coverage, UX, onboarding, analytics and self-serve flows

Infra, Legal & Ops (20%)

Server costs, tools, patent prosecution, basic operations

Key milestones for this round:

Launch paid tiers (Month 1)

Reach \$5K MRR in 90 days (~100 paying users at \$49.99/mo)

Establish at least one repeatable acquisition channel with positive unit economics

Maintain 6+ months runway at the end of this period to set up the next round

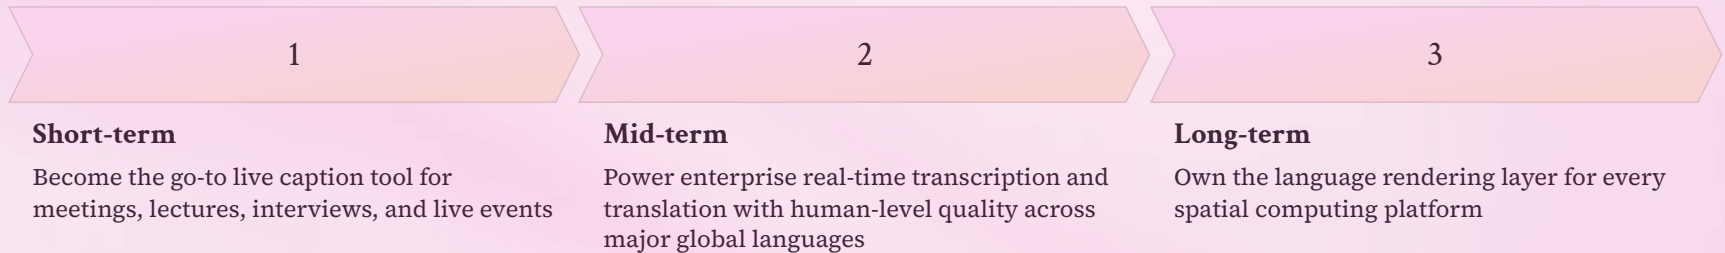
Investment Opportunity

Back a deep-tech voice infrastructure company at the moment it turns validated usage into revenue. Early allocation at \$10M cap for the first \$100K, with the broader round anchored around a \$10–12M cap.

Vision & Contact

LansonAI is building the Voice Context Layer between people and machines — starting with Lanson Live, which makes live speech easy to read and follow through stable real-time captions.

Roadmap



Talk is cheap. Try it yourself.

Lanson Live

live.lansonai.com

no login required, 5 min session limited

Video: Story of Serenity

www.lansonai.com

Email

zhen@lansonai.com

Appendix A: Methodology & Data Sources

All performance claims are based on internal measurements and publicly available data. Full methodology available on request.

Batch Processing Speed

120×*

Measured end-to-end: 1 hour of audio processed in under 30 seconds using Trinity Engine's serverless parallel architecture. Benchmark conducted on internal test suite. Actual performance may vary by audio length.

Zero-Edit Rate

90%*

Defined as: output requiring no human correction before publication. Measured across internal test sessions during beta. Sample size and methodology available on request.

Cost Reduction

80%*

Based on internal production cost benchmarks (measured) compared to publicly listed per-minute pricing from major providers including AssemblyAI, Deepgram, Comparison reflects cost-to-deliver, not retail pricing.

* These figures reflect internal measurements and/or publicly available market data. LansonAI welcomes technical due diligence at any level of scrutiny.