



Communication Intelligence & Cognitive Profiling

We **don't** just **transcribe** what people say.
We **decode** how they **think**, **listen**, and process information
- **synchronizing communication across humans and AI.**




Vitalii Holiienko
Chief Executive Officer



Serhii Kramer
Chief Technology Officer



Dr. Sergiy Shyshko
Chief AI Scientist



Beyond Transcripts: The Unmapped 80% of Voice Data

Traditional systems focus on what was said. Modern communication requires understanding how people think, process, and engage.

The Problem (Traditional Text-Centric Systems)	The Solution (SAPYSYNC Cognitive Approach)
Context Flattening: Transcripts preserve raw words but eliminate response latencies, speech rhythm, pauses, and turn-taking balance.	Signal-Centric Intelligence: Extracts speech timing, tempo shifts, and engagement patterns directly from acoustic streams.
No Behavioral Depth: Standard STT (Speech-to-Text) tools fail to capture hesitation, cognitive overload, or active listening dynamics.	Cognitive & Behavioral Profiling: Measures attention, comprehension, and speaker dynamics in near real-time.
Heavy Data & Privacy Overhead: Storing full conversation text creates massive compliance risks (GDPR, HIPAA, AI Act) and high storage costs.	Privacy-First Standard: Processes rich interaction metrics without mandatory storage of full transcript content.

Dual-Engine Architecture: Speed Meets Deep Intelligence

A pragmatic processing model combining near real-time streaming with asynchronous deep analytics

Engine 1

Fast Engine (Near Real-Time Streaming):

- **Low-Latency Windowing (~2–3s):** Replaces false "instantaneous" promises with realistic sub-3-second audio buffering to capture live speech dynamics.
- **Live Operational Nudges:** Delivers instant feedback during calls—detecting active listening drops, friction, and pacing mismatches.
- **Sub-Second Signal Extraction:** Tracks micro-pauses and response latency without wait times.

Engine 2

Deep Engine (Async Post-Processing):

- **Historical Archive Profiling:** Runs heavy neural models on call recordings and historic databases to build comprehensive speaker profiles.
- **Long-Term Cognitive Baselines:** Analyzes behavioral trends, emotional resilience, and interaction consistency over time.
- **Standardized Reporting:** Generates compliance-ready interaction logs, clinical documentation, and strategic coaching insights.

Predictive Speaker Profiling & Smart Matching

Transforming historical conversation archives into persistent cognitive profiles and optimal interaction strategies.

Archive Processing → Cognitive Profile → Predictive Match / Pre-Call Nudge

Asynchronous Speaker Profiling:

- **Persistent Cognitive Profiles:** Processes historical call records to establish unique behavioral baselines for each participant-mapping baseline pacing, response latency, and hesitation patterns.
- **Communication Persona Identification:** Categorizes speakers by processing style (e.g., highly analytical with long latencies vs. high-tempo dynamic decision-makers).

Predictive Matching & Routing:

- **Smart Operator/AI Matching:** Routes incoming callers to the human agent or Voice AI persona with the most compatible communication dynamics to maximize rapport.
- **Zero-Latency Pre-Call Strategy:** Delivers tailored interaction prompts (e.g., "Maintain a 1.5s response pause") at the onset of the call-eliminating real-time computation bottlenecks.

Modular Intelligence: Vertical "Specialist" Modules

A universal core powered by plug-and-play domain ML engines tailored to industry-specific behavioral needs.

Domain Specialist Add-ons (Selectable Engines):

- **Sales & Negotiation Specialist:**

Identifies buying signals, price-hesitation latencies, objection patterns, and closing readiness.

- **Healthcare & Therapy Specialist (Mental Health):**

Tracks acoustic biomarkers of depression, anxiety, and cognitive fatigue analysing speech rhythm and pause distribution over time.

- **Logopedics & Speech Therapy Specialist:**

Provides automated speech fluency analytics, tracking articulation timing, stuttering patterns, and speech recovery progress.

- **Emergency & Public Safety Specialist (911 / 112):**

Measures caller panic levels and cognitive overload in real time to guide dispatchers during crisis de-escalation.

- **EdTech & Training Specialist (Comprehension & Attention):**

Evaluates active listening, student engagement, and topic comprehension by detecting hesitation frequency and pacing mismatches.

Privacy-First Architecture & Regulatory Compliance

Building an ethical, enterprise-ready communication layer compliant with strict global data regulations.

Privacy by Design (Technical Security):

- **No Mandatory Text Storage:** Processes acoustic and behavioral metrics (timing, rhythm, latency) directly from audio streams without storing sensitive text transcripts.
- **Zero PII Exposure:** Automatically decouples personal identifiable information (PII) from speech metrics, protecting individual identity.
- **On-Premise & Isolated Deployment:** Flexible enterprise deployment options—cloud, hybrid, or fully on-premise for high-security environments

Global Regulatory Readiness:

- **EU AI Act & GDPR Compliant:** Fully aligned with European Union guidelines for biometric, AI, and data processing safety.
- **HIPAA Ready for Healthcare:** Meets healthcare privacy standards for medical, therapy, and speech analytics applications.
- **Ethical AI Alignment:** Eliminates invasive "lie detection" pseudoscience in favor of objective, scientific cognitive friction metrics.

Business Model & Go-To-Market Strategy

High-margin recurring SaaS model driven by direct pilot deployments and value-based subscription tiers.

Unit Economics & Pricing Tiers:

- **Flexible SaaS Tiers:**
 - **Standard Async Analytics: €0.015 / min** (Bulk processing of uploaded call archives and historical audio)
 - **Live/Near Real-Time Assistant: €0.03 / min** (Streamlined processing via Web Interface / WebRTC)
 - **Specialist Engine Add-ons:** Flat monthly subscription per specialized module (Health, Emergency, Sales)
- **Unit Economics (Per 60k Minutes/Mo):**
 - **Monthly Revenue: ~€1,800 (~€21.6K ARR per mid-size client)**
 - **Infrastructure Cost: ~€140 / month**
 - **Estimated Gross Margin: ~92%**

Direct Go-To-Market & Pilot Strategy:

- **Founder-Led Direct Sales:** Direct outreach targeting BPO contact centers, psychological practices, and specialized dispatch teams for paid pilot programs.
- **Turnkey Web Interface:** Frictionless onboarding- clients upload audio records or use our web app with zero complex technical integration required.
- **Value-Based Retention:** Converting successful pilot projects into annual enterprise SaaS contracts with predictable monthly volume commitments.

Competitive Landscape: Redefining Voice Intelligence

Moving beyond basic text transcription and generic sentiment analysis toward privacy-first cognitive profiling.

Feature / Capability	Traditional STT (Gong, Otter, Fireflies)	Emotion AI (Voice Stress Analysis)	Enterprise LLMs (Custom Audio GPTs)	SAPYSYNC Platform
Primary Focus	Raw text & keywords	Pitch & basic emotional stress	Full transcript reasoning	Cognitive & Behavioral Profiling
Acoustic Timing Analysis	✗ Text-only (ignores timing)	⚠ Limited micro-tremor	✗ Ignored in text pipeline	✓ Full Latency, Rhythm & Pacing Map
Privacy & GDPR Compliance	⚠ Stores full transcripts	✓ Low text storage	✗ High risk (Full data retention)	✓ Privacy-First (No-Text Streaming)
Modular Domain "Specialists"	✗ One-size-fits-all	✗ Niche-specific	⚠ Requires heavy prompt engineering	✓ Plug-and-Play Vertical Engines
Infrastructure & Compute Cost	Medium	Low	Very High (Expensive GPU load)	✓ Low Cost (~92% Gross Margin)

Key Strategic Moat (Unfair Advantage):

- The "How" Over "What": While competitors burn compute power transcribing and parsing words, SAPYSYNC extracts actionable intelligence directly from acoustic timing, hesitation, and active listening dynamics.
- Compliance-First Architecture: Complete freedom from transcript leaks and zero compliance hurdles under EU AI Act & HIPAA standards.

Product Roadmap: Strategic Execution Phases

Phased deployment path from core web-based analytics to live streaming assistants and specialized vertical engines.

Phase 1 — September 2026: Core SaaS & Nextcloud Integration (MVP 1)

- **SaaS Upload & Post-Call Deep Engine:** Direct file upload interface for recorded calls and bulk audio archives.
- **Native Nextcloud Talk Webhook:** Automated post-call processing via Nextcloud's open-source privacy-first ecosystem.
- **Core Cognitive Dashboard:** Baseline reporting on speech timing, response latencies, pause distribution, and interaction friction.

Phase 2 — January 2027: Cross-Platform Streaming & Near Real-Time (MVP 2)

- **Universal Meeting Bots:** Rapid deployment across Zoom, MS Teams, and Google Meet via WebRTC meeting bot architecture (Recall.ai/Playwright).
- **Low-Latency Streaming (~2–3s):** Real-time audio windowing for live conversation guidance and pacing analysis.
- **Live Operational Nudges:** Instant feedback triggers during live calls to prevent active listening drops and communication friction.

Phase 3 — April 2027: Specialized Vertical Engines & Smart Matching

- **Modular "Specialists":** Deployment of plug-and-play domain engines (Sales, Healthcare/Therapy, Speech Correction, Emergency Dispatch).
- **Predictive Speaker Matching:** Historical profile processing to route incoming callers to compatible operators or Voice AI personas.
- **Enterprise Security & Compliance:** On-premise deployment packages, HIPAA validation, and EU AI Act regulatory compliance.

Investment Ask: Funding the Next Phase of Growth

Raising Pre-Seed capital to complete MVP development, validate direct market demand, and achieve core tech milestones.

The Ask & Round Structure:

- **Raising: €350,000** (Pre-Seed Round via SAFE / Convertible Note)
- **Runway Target: 18 Months** of operational runway
- Core Milestones to Unlock:
 - Launch MVP 1 (SaaS Upload & Post-Call Deep Engine)
 - Onboard 10–15 Paid BPO & Healthcare Pilot Clients
 - Validate unit economics and prepare architecture for Near Real-Time Assistant (MVP 2)

Strategic Use of Funds:

- **55% — Core R&D & Engineering:** Building the acoustic timing pipeline, post-call analytics engine, and modular domain algorithms.
- **25% — Go-To-Market & Pilot Acquisition:** Founder-led sales execution, marketing, legal support, and onboarding setup for B2B pilots.
- **20% — Cloud Infrastructure & Operations:** Server compute costs, GPU processing, security architecture, and regulatory compliance (GDPR/HIPAA).

Founding Team



Vitalii Holiienko

Chief Executive Officer



15+ years in Strategy, Sales & Business Development
Driving product vision, partnerships and go-to-market strategy.



Serhii Kramer

Chief Technology Officer



30+ years in Software Development, R&D and AI.
Expert in scaling complex technical platforms and implementing cutting-edge AI solutions.



Dr. Sergiy Shyshko

Chief AI Scientist



PhD in Applied AI & Complex Systems
Leading AI research, machine learning and behavioral speech modeling.



THANK YOU

**SAPYSYNC Restoring alignment in communication
between people, systems and AI**

Understanding creates trust

Trust creates better communication

Let's build the future of AI communication together

Get in touch:

[LinkedIn](#) | [Mail](#) | [WhatsApp](#)

