



LONG-FORM VISION · FOR INVESTORS · 2026

# The wisdom substrate. Willingly given. Permanently kept.

Frontier AI is starved of one specific data class: lived human judgment. It is not on the internet — it lives in seniors' minds. Sacred Blue is building the substrate that captures it on the way out, and the model that delivers it back to the people who need it.




# The data frontier AI cannot scrape.

Every frontier model is trained on the indexed internet. That corpus is enormous — and almost entirely missing one thing: the lived moral judgment of people who have actually navigated love, grief, failure, illness, leadership, and death.

No amount of scraping will fix that. The wisdom of a generation that didn't grow up online has never been written down. It exists in conversation, in memory, in the small lessons a grandmother says once and never repeats.

## THE BET

**A model trained on willingly-given senior wisdom will reason about human situations in a way no general-purpose LLM can.**



Lived moral judgment  
Not online. In memories.



THE STUDIO

# Sacred Blue.

Sacred Blue is the studio behind Waterfall, River, and the Wisdom Model. We build on top of human wisdom — we don't scrape it. Every product is downstream of the same conviction: the substrate that makes AI useful for human life is the part of human life that AI cannot generate on its own.

## OPERATING PRINCIPLES

- Consent is the product, not a setting. Every grant is per-story, per-purpose, revocable end-to-end.
- Provenance is permanent. Every training pair points back to the human who taught it.
- Distribution follows trust. Facilities, faith communities, families — not feeds.
- The studio compounds. Every product feeds the substrate that makes the next one better.

Products: Wisdom Model API · Water · River

Models: fine-tunes on the Waterfall corpus

Substrate: consented, attributable senior wisdom

A studio aligned vertically: one substrate, one model layer, many products.

# A closing window meets an open frontier.

Two curves are converging. The largest senior generation in history is at the end of life — and the AI industry, having exhausted the public web, is hunting for the next training-data substrate. The right window to build this is now.

**2.1B**

People over 60 by 2050. Doubling from ~1.1B today.

UN WPP 2022

**10K**

US seniors passing each day — each carrying a lifetime out the door.

CDC NVSR 2024

**~0%**

Share of frontier-LLM training data that is consented, attributable senior wisdom.


ESTIMATE

**\$500B**

Projected AI-agent spend by 2030 — every agent will need grounded judgment.

GARTNER / MCKINSEY

# Capture on one side. Inference on the other.



## WATERFALL · THE GIVING SIDE

### Seniors deposit stories.

Voice, typing, photos of letters and journals. Multi-tier consent, jurisdiction-aware, revocable. Stories never enter the corpus without an explicit grant from the person who told them.

## RIVER · THE RECEIVING SIDE

### Youth withdraw answers.

A teenager asks about grief, failure, identity, faith. River retrieves real elder recordings that answer it — with attribution back to the senior who taught it. Inference is grounded in the substrate.

# A capture apparatus built for trust.

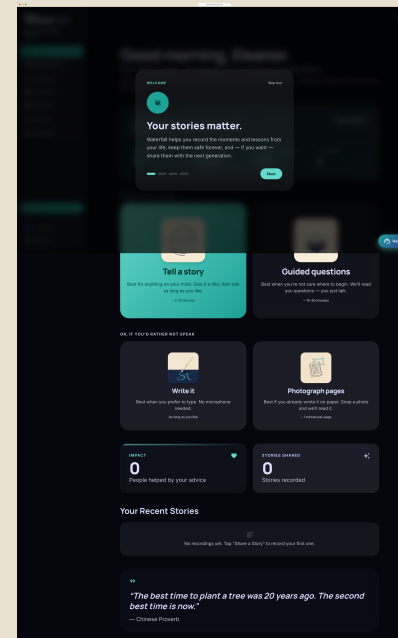
Waterfall meets seniors where they are. Voice for those who prefer to speak. Typing for those who write. Photographs for handwritten letters and diaries. Three modes converge into one editable transcript per story.

## CHANNELS

- Direct app — senior records on their own device.
- Facility partner — staff-assisted sessions in LTC, retirement, hospital volunteer.
- Family-mediated — youth or adult child captures on a parent's behalf.
- Wisdom Circle — group sessions for shared lived experience.

## CONSENT MODEL

Four tiers, each individually revocable: platform-only 'model training 'API access 'attributed publication. Default is the most restrictive. Revocation scrubs every downstream artifact within 30 days.



# Specialised questions, designed by research.

A senior who is asked "tell me about your life" delivers a polished autobiography. A senior asked the right question, in the right order, surfaces the moment they actually learned the lesson — with the principle, the why, and the application intact.

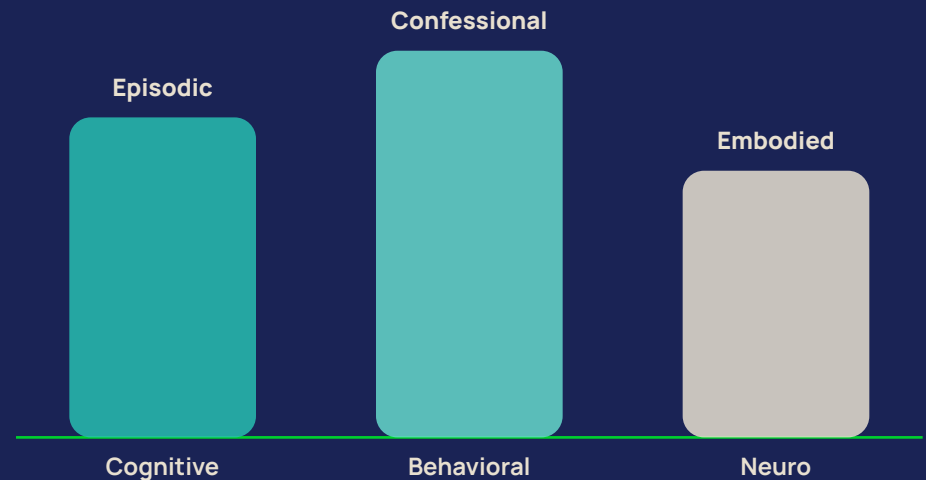
## THE CORPUS

- 100 questions across 5 wisdom domains, 26 sub-domains.
- 8 follow-up probes calibrated to elicit principle, then why, then application.
- 19 session templates — facility-length (30–45 min), app-length (10–15 min), Wisdom Circles for shared experience.

## THE PROTOCOL

### Story 'Principle 'Why 'Application

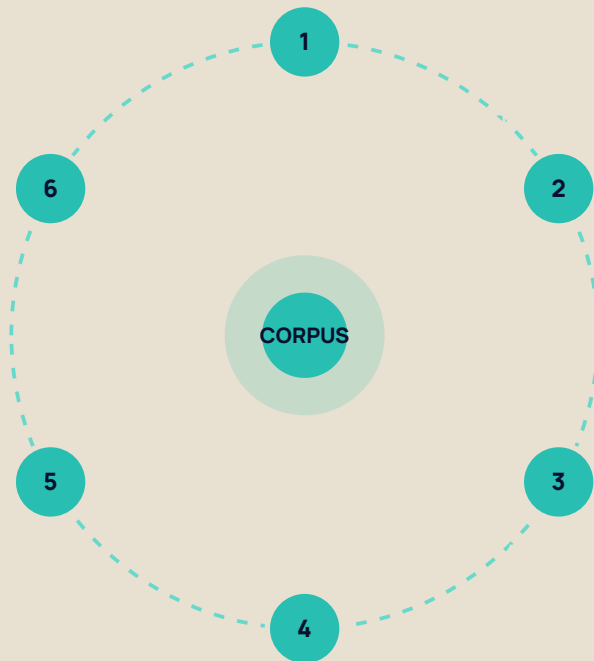
Every session walks the senior down the same ladder, in their own words — the structure that turns a memory into trainable judgment.



## THREE SCIENTIFIC PILLARS

- Cognitive — episodic activation (Tulving), semantic priming, bridge questions.
- Behavioral — confession framing, third-person projection, temporal distancing.
- Neuroscientific — narrative prompts (default mode network), analytical prompts (PFC), embodied prompts (somatic memory).

# From one voice to a trainable corpus.



Six stages, each a separately-deployed Pub/Sub worker, each idempotent. Audio enters as raw consented speech and leaves as a structured WisdomRecord with quality signals, emotion profile, moral dimensions (Haidt), and provenance attached.

## STAGES

- Transcribe — Google STT with diarization.
- Annotate — Vertex AI Gemini extracts principles + moral foundations; Hume.ai derives emotion.
- Validate — automatic quality gates plus human review queue.
- Embed — Vertex embeddings into Vector Search.
- Train — generate training pairs (situationresponse, dilemma'-reasoning, emotionscenario), export to BigQuery + JSONL for Vertex fine-tuning.

Revocation is end-to-end: pulling a consent grant scrubs the record, the audio, the vector, and every training pair derived from it — within 30 days, with audit log.

# Where the substrate becomes counsel.

Inference happens on the River side. A young person asks a question — a real one, the kind they wouldn't Google. River answers two ways:

## 1. Real recordings, attributed.

Cluster-matched elder recordings play directly. The senior is named only if they granted attribution; otherwise their voice is anonymised but the lesson is intact.

## 2. Wisdom Model synthesis.

When no exact match exists, a Vertex fine-tune trained on the corpus reasons in the seniors' register — cautious, plural, never pretending to be any one of them.

### WHAT RIVER REFUSES TO BE

- Not a chatbot impersonating any one senior.
- Not a feed. No likes, no engagement loops.
- Not therapy — stories, not clinical advice.
- Not polished content — real, sometimes uneven.

### THE ARC

Each elder lesson on River permanently strengthens the model. Each youth question reveals what the substrate is missing — closing the loop back to Waterfall's next collection cycle.

# Why a scraper cannot replicate this.

The moat is not the model. It is the substrate — the consented, attributable, emotion-tagged corpus underneath. Replicating it requires a relationship with the storyteller. That relationship is the asset.

## WHAT COMPETITORS CANNOT COPY

- Provenance — every record traces back to a person and a grant.
- Consent posture — four tiers, revocable, jurisdiction-aware (HIPAA, BIPA, MHMDA, PIPEDA, Quebec Law 25, PHIPA, HIA, PIPA, PHIA).
- Emotion grounding — every clip Hume-tagged where lawful.
- Multimodality — audio, transcript, photos, journal scans — against one canonical record.
- Distribution — facility partnerships, family-led capture, trust-first not feed-first.



A model trained on the open web ingests anyone. A model trained on the substrate ingests no one without their permission — and that is precisely why it can be trusted in the room.

# Two products. One foundation.

## WISDOM MODEL

### Moral reasoning, on demand.

A Vertex fine-tune trained on the corpus. Any AI agent can call it for moral reasoning, life guidance, and human judgment. Three usage tiers: Explorer (free, 100 q/mo), Builder (\$0.01/query, metered), Enterprise (custom, BAA).

- situation\_response training pairs
- dilemma\_reasoning training pairs
- emotional\_scenario training pairs

## WATER

### Emotional intelligence engine.

A separate API trained on the emotion-tagged slice of the corpus. Three endpoints: Read (read the room), Respond (suggest an emotionally intelligent response), Calibrate (track an emotional trajectory across a conversation).

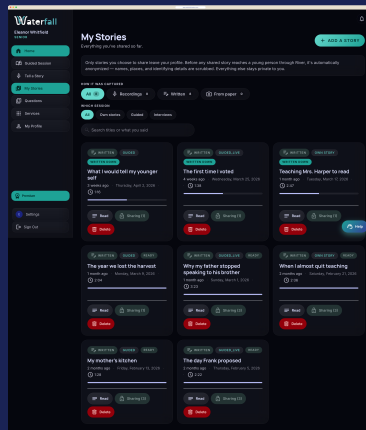
- WaterReadRequest 'EmotionalRead
- WaterRespondRequest 'WaterResponse
- WaterCalibrateRequest 'WaterCalibration

Both APIs are downstream of the same WisdomRecord schema. Adding a senior to Waterfall today permanently improves both products tomorrow.


BUILT

# The substrate is not theoretical.

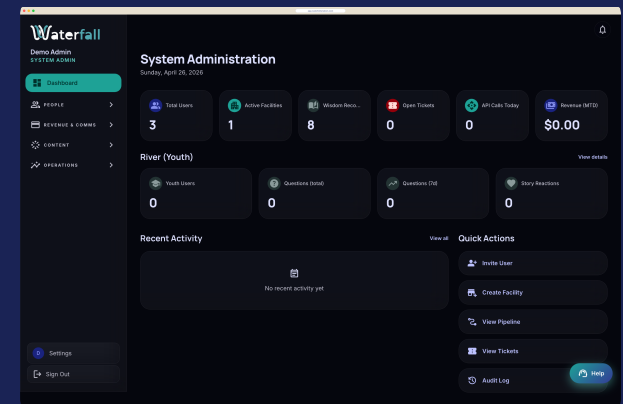
The Waterfall platform is shipped: senior collection, facility administration, system observability, the full Pub/Sub processing pipeline, the Wisdom API with Stripe billing, ten-language marketing site, and a React Native mobile app.



Senior — session selection



Facility — session monitoring



System admin — platform overview

4

apps in monorepo  
(web, mobile, marketing,  
screenshots)

7

pipeline workers (transcribe ' training)

10

languages live on the  
marketing site

100

research-grounded questions  
in the framework

# The corpus compounds.

---



---

<b>Year 1</b>	<b>Pilot</b>	5 facilities · ~750 stories · Wisdom Model v1 · IL/WA/CA/QC compliance live · SOC 2 Type II.
<b>Year 3</b>	<b>Substrate</b>	500 facilities · 250k stories · 25 languages · Wisdom API GA · River public launch · HITRUST e1.
<b>Year 5</b>	<b>Default layer</b>	Wisdom Model is the standard substrate any AI agent calls for human judgment. River is how a generation talks to its elders.

---

# Help us close the window.

## USE OF CAPITAL

- Collection capacity — facility partnerships, staff training, family-led expansion.
- Substrate quality — question framework R&D, multilingual coverage, validation review.
- Model layer — Vertex fine-tunes on the pilot corpus; eval harness vs. general-purpose baselines.
- Compliance — SOC 2 Type II, HITRUST e1, healthtech counsel, multi-state consent posture.
- Sacred Blue R&D — the next products downstream of the substrate.

The unit economics are tight: variable cost per 10-minute story is ~\$0.10. Capital is for collection volume, not infra.

## CONTACT


### Maurice Savery

Founder, Sacred Blue · Waterfall

[msavery@gmail.com](mailto:msavery@gmail.com)

[waterfallwisdom.com](http://waterfallwisdom.com)

**Stories worth saving.  
Permanently kept.**



A Sacred Blue studio