



CAPABILITIES

Winsen Labs

An applied AI lab.

Financial Services · Manufacturing · Logistics · Healthcare

winsenlabs.com

WHY NOW

AI can do far more than most organisations use it for.

Models now read documents, follow a process and take actions.
They get noticeably better every few months.

Most organisations are still at the pilot stage: a few demos and internal tools, and very little running at volume.

We build the systems that put this capability into everyday work.

We take what new models can do and turn it into systems people use every day.

RESEARCH LAB

Builds the models.

Advances what AI can do and publishes the results.
Putting it to work is left to others.

SOFTWARE COMPANY

Sells a finished tool.

Designed around what AI could do when it was built,
and updated on a release cycle.

APPLIED AI LAB

Builds working systems on the latest models.

We test each new capability against real work, build what holds up, and update our systems as models improve.

HOW THE LAB RUNS

01

Explore

Test new models on real work.

02

Build

Build what works into products.

03

Open source

Release useful tools as open source.

04

Teach

Teach the methods at Winsen Academy.

What we have built so far.

PRODUCTS

Winsen One

AI employees for financial services, starting with risk.

Winsen Rho

Monitoring and governance for every AI agent an institution runs.

FOUNDATIONS

Alfred

The in-house harness every AI employee runs on.

Pioli

Physical observability for plants, ports and cities, built on Winsen Ontology.

IN THE OPEN

OCSO

Open customer success orchestration

OVO

Open voice orchestrator

AI employees for risk.

Each AI employee has one role, one queue of work and a defined output. It reviews every item in the queue rather than a sample, runs inside the institution's environment, and its work only counts once a named person signs it.

CITES SOURCES

Every finding resolves to the document, the clause and the page.

STAYS INSIDE

Documents and findings never leave. Inference runs in India.

A PERSON SIGNS

It never reviews its own work. Every action is logged and replayable.

TEN SEATS UNDER A CRO

Credit risk review DEPLOYABLE	Operational risk DEPLOYABLE
Compliance and regulatory change IN PROGRESS	Liquidity and ALM risk IN PROGRESS
Market risk IN PROGRESS	Fraud risk IN PROGRESS
Information and cyber risk IN PROGRESS	Model risk IN PROGRESS
Enterprise and integrated risk IN PROGRESS	The manager IN PROGRESS

Every AI employee runs on Alfred, our in-house harness.

Alfred gives each employee a persistent machine of its own, deployed inside the institution’s VPC. Its reasoning, credentials, memory, work state and data all stay on infrastructure the institution controls.

WHAT AN EMPLOYEE GETS


Its own machine	Persistent, so work carries on across days rather than resetting each session.
Memory and work state	Responsibilities, actions and outcomes recorded in a local database.
A credential vault	Keys are encrypted locally and can be rotated without a redeploy.
Computer use, sandboxed	The VPC is the sandbox. The employee can execute freely inside it, and nowhere else.

WHAT THE INSTITUTION KEEPS

Approval before effect	Each approval is bound to the exact proposal, checked again before anything happens.
Guarded outbound calls	Every outbound request is checked by address before it leaves, and after every redirect.
Its choice of model	The model route is set explicitly and recorded on every run, never guessed.
Independence	The cloud panel is optional. Local work does not depend on it.


An AI governance platform: a record of every AI decision, checked against what actually happened.

Rho tracks every AI agent an institution runs, whether built in-house or bought from a vendor. It compares each decision with what happened next in the institution's own systems, and turns that record into reports for the board and evidence for the regulator.



Pioli gives AI a live view of plants, ports and cities.

Pioli is our physical observability platform. It records what is happening on site as structured data, which AI needs before it can take any action in the physical world.



Most of what we build along the way is released as open source.

OPEN SOURCE

OCSO

Open Customer Success Orchestration

One open layer where AI agents and people serve customers on WhatsApp, web chat, Slack and Microsoft Teams, with human handoff both ways, maker-checker approval on every change and a tamper-evident audit trail. Self-hosted, Apache 2.0, pre-1.0.

ocso.winsenlabs.dev

IMPLEMENTATION AVAILABLE

OVO

Open Voice Orchestrator

Build, run and observe voice agents on your own infrastructure, where everything is a plugin: the conversation engine, speech, telephony, tools and the console. Runnable today, with production certification pending.

github.com/winsenlabs/ovo

One open layer where AI agents and people serve customers together.

Customer service is usually spread across channels, tools and teams, with AI added at the edges. OCSO puts it in one place, on the organisation's own infrastructure.

CHANNELS

WhatsApp, web chat, Slack, Teams

Agents and people answer customers wherever they already are.

HANDOFF

Human handoff, both ways

A person can step in at any point and hand back when done.

TOOLS

Your own systems over MCP

Agents act through the tools the organisation already runs.

CONTROL

Maker-checker on every change

No change goes live without a second person approving it.

RECORD

Tamper-evident audit trail

Every conversation and action can be reviewed later.

OPEN

Self-hosted, Apache 2.0

Pre-1.0 and under active development.

One government service desk. *A citizen explains their need once.*

A citizen reaches out directly by voice, WhatsApp or app, or walks into a Common Service Centre where a Village Level Entrepreneur guides them with an AI assistant. Both paths meet in one service desk across departments. Officers decide, in their own systems.

OCSO runs the service desk: one case per citizen, followed across every department until it is resolved.

WHAT OCSO ALREADY DOES

WHAT THE DESK NEEDS

AI agents and people on WhatsApp, web chat and more	→	Citizens reach out on the channel they already use
Human handoff, both ways	→	The VLE takes over when the citizen needs a person
Tools from your own systems over MCP	→	One case across Revenue, Agriculture, Health, Pensions
Maker-checker approval on every change	→	Officers decide; AI assists, people are responsible
Tamper-evident audit trail	→	Every step recorded, tracked to resolution

Own your voice.

OVO is our open voice orchestrator. Organisations build, run and observe voice agents on their own infrastructure, and choose the provider behind every part of the call.

EVERYTHING IS A PLUGIN

Conversation	Scripts, FAQ matching, supplied-context dialogue, tool-using agents
Voice	Speech recognition, synthesis, turn detection, playback
Connectivity	Telephony, media transports, business integrations
Operations	Recordings, secrets, telemetry, evaluations, cost accounting

Active calls pin their configuration and plugin versions.

FOUR WAYS TO BUILD AN AGENT

Message with variables

Read an appointment reminder or account update.

NO LLM

Script and FAQ

Follow an approved script and answer known questions.

NO LLM

Supplied-context conversation

Answer questions using only the information given to the agent.

LLM, BOUNDED

Agent with tools

Check availability or update a business system.

LLM + PERMITTED TOOLS

We start with your problem, not our product.

01 · CONSULTING

Explore

Bring us a process that has been too slow, costly or difficult to improve. We assess what current AI can do for it and tell you whether it is worth building.

02 · FORWARD-DEPLOYED

Build

Our engineers work with your team from prototype to production, around your systems and constraints, and hand over to the people who will run it.

03 · ACADEMY

Learn

Winsen Academy teaches people to work with AI from first principles, taught by the engineers who build with it. The first courses are being written now.

Reimagine and Rebuild.

Work is better with Winsen.