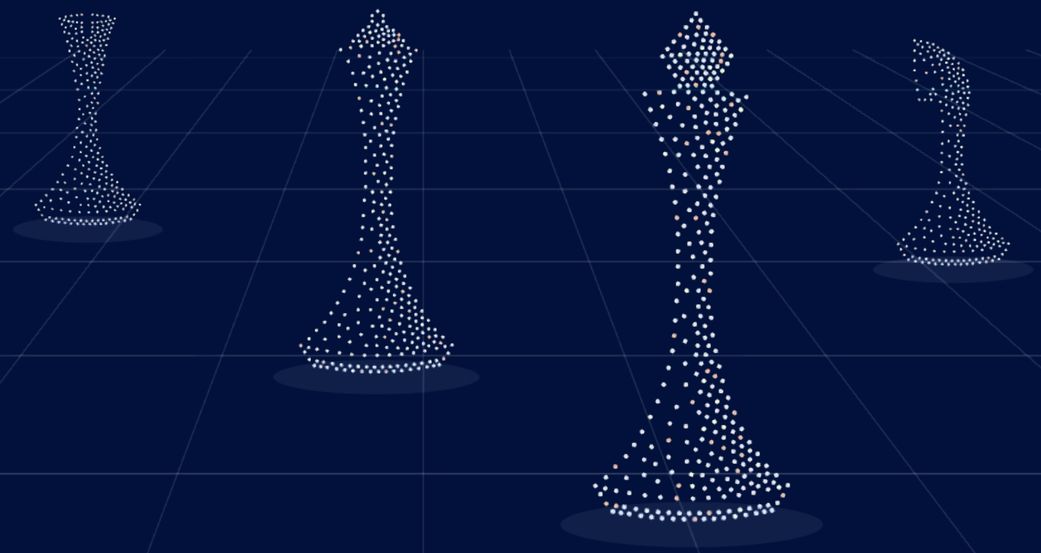
PITCH DECK · NOVEMBER 2026

AlphaEdge

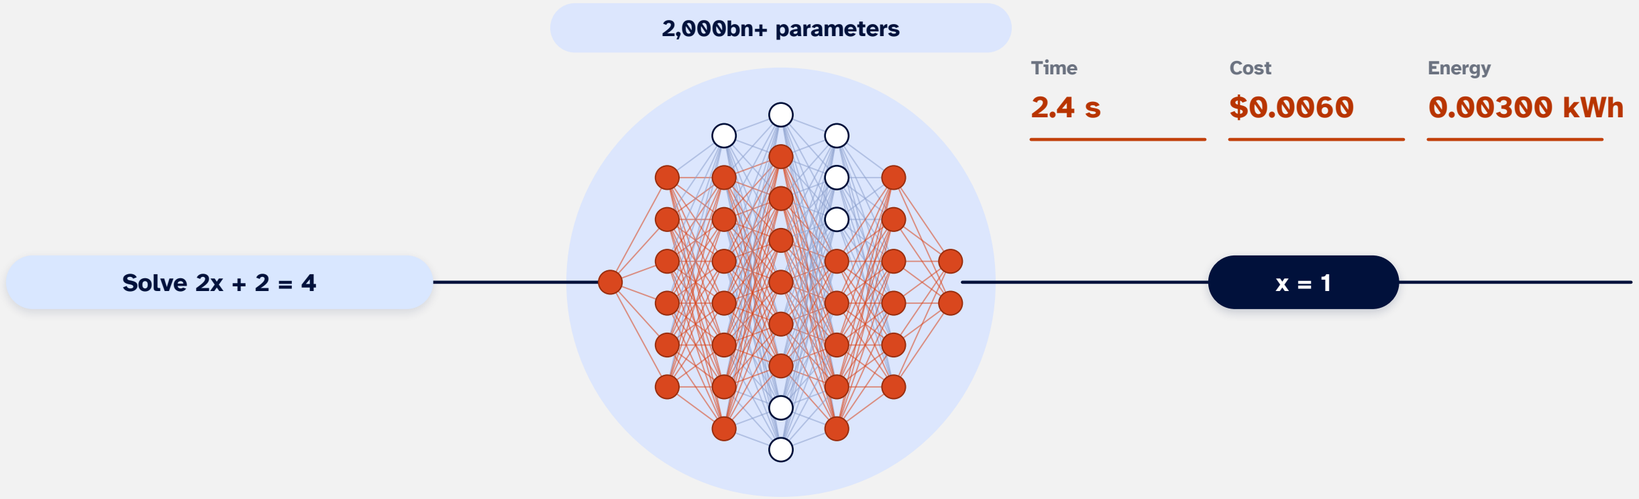
Building AI where it matters



Efficient Language Models · Specialized, efficient and sovereign AI

The limits of a one-size-fits-all model

LLMs are **oversized and generalist**: the simplest question mobilizes the knowledge of all humanity

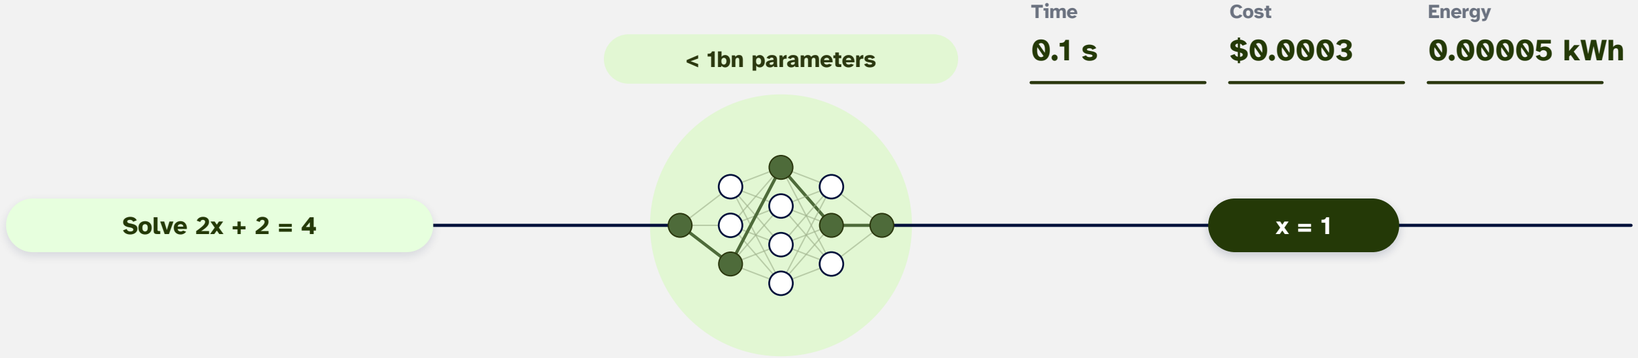


COST	SOVEREIGNTY	ENERGY	LATENCY
$12 / M tokens	**Extraterritorial**	**High**	**> 2 s, delayed**

AlphaEdge

Our solution: efficient, specialized models

For 90% of use cases, a **well-optimized lightweight model** performs better and costs less. We call our model families **ELM (Efficient Language Models)**.

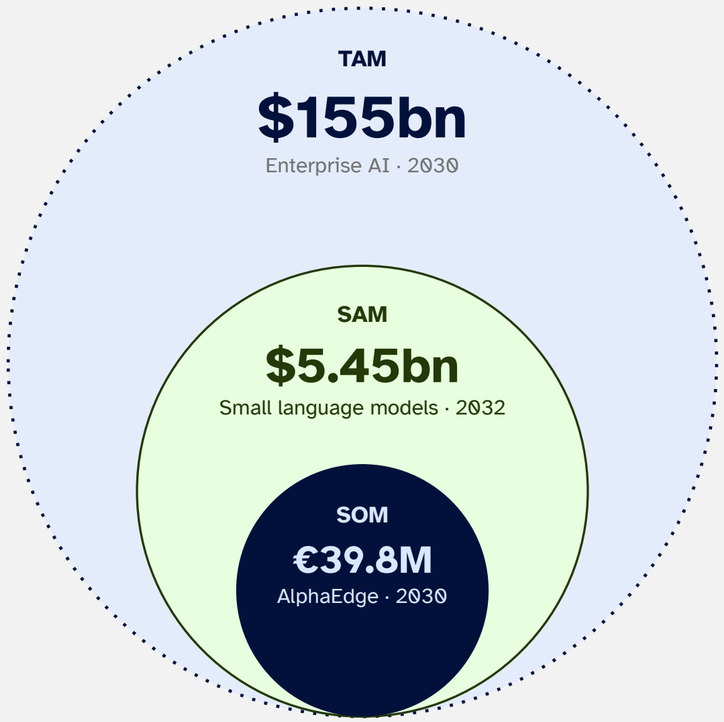


COST	SOVEREIGNTY	ENERGY	LATENCY
$0.5 / M tokens	**Full**	**Low**	**0.1 s, real time**

AlphaEdge

The **B2B** AI market: **over $155bn in 2030**

Shift to specialized models: by 2027, enterprises will use small task-specific models 3× more than general-purpose LLMs (Gartner).



TAM · ENTERPRISE AI MARKET

$155.2bn in 2030

+37.6% per year (Grand View Research)

SAM · SMALL LANGUAGE MODELS

$5.45bn in 2032

+28.7% per year, the segment of our ELMs (MarketsandMarkets)

SOM · OUR BUSINESS PLAN (BASE CASE)

2027 · FRANCE FOCUS

€1.1M revenue

20 new customers · €274k ARR

2030 · FRANCE + 2 COUNTRIES

€39.8M revenue

342 active contracts · 17% international

Our 3 priority target sectors

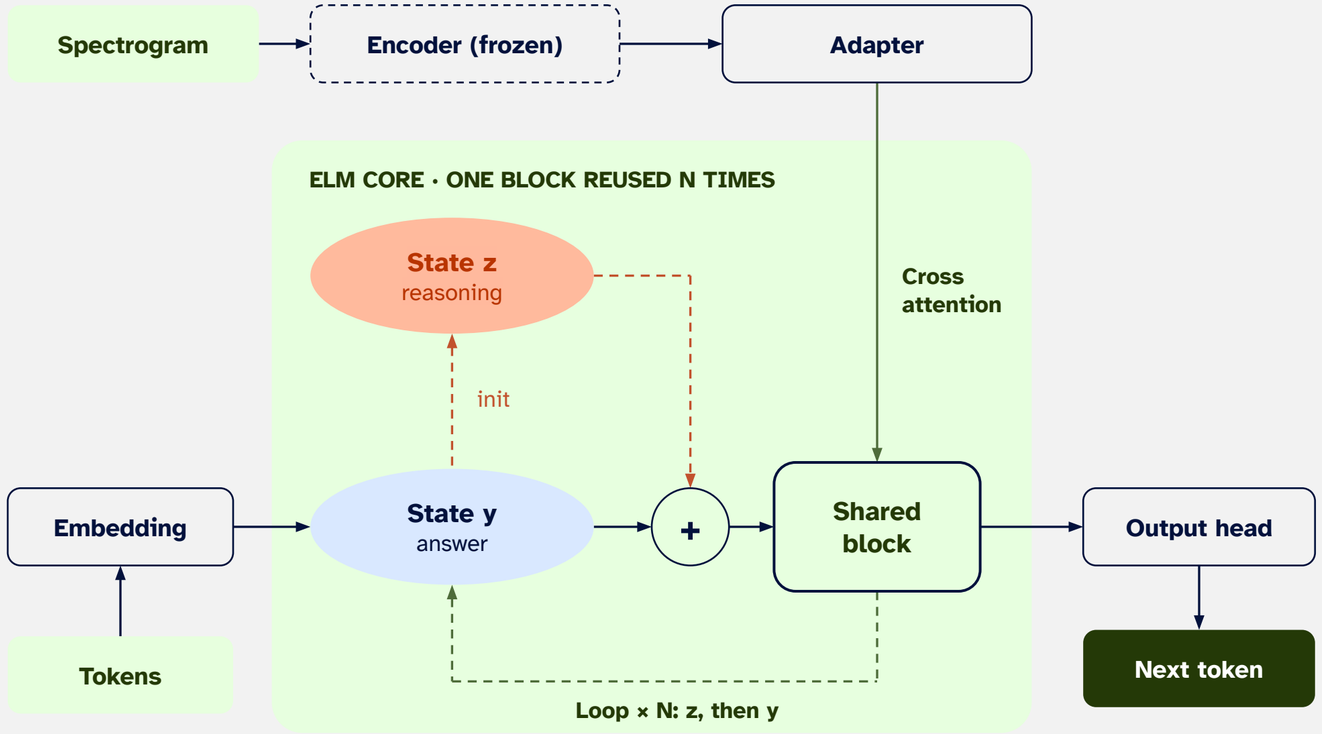
42%	31%	27%
Industry & manufacturing	**IT services & integrators**	**Banking, finance & insurance**

AlphaEdge

Our lab designs efficient AI models

AlphaEdge designs the architecture, then trains it on domain data: a **native ELM model,** 100% proprietary.

EXAMPLE · NATIVE ELM ARCHITECTURE (AUDIO → TEXT)



01 · RESEARCH & ARCHITECTURE

Inventing new architectures

Our lab designs its own mathematical model building blocks, built from day one to be frugal and high-performing.

02 · DOMAIN DATA

Specializing each model through data

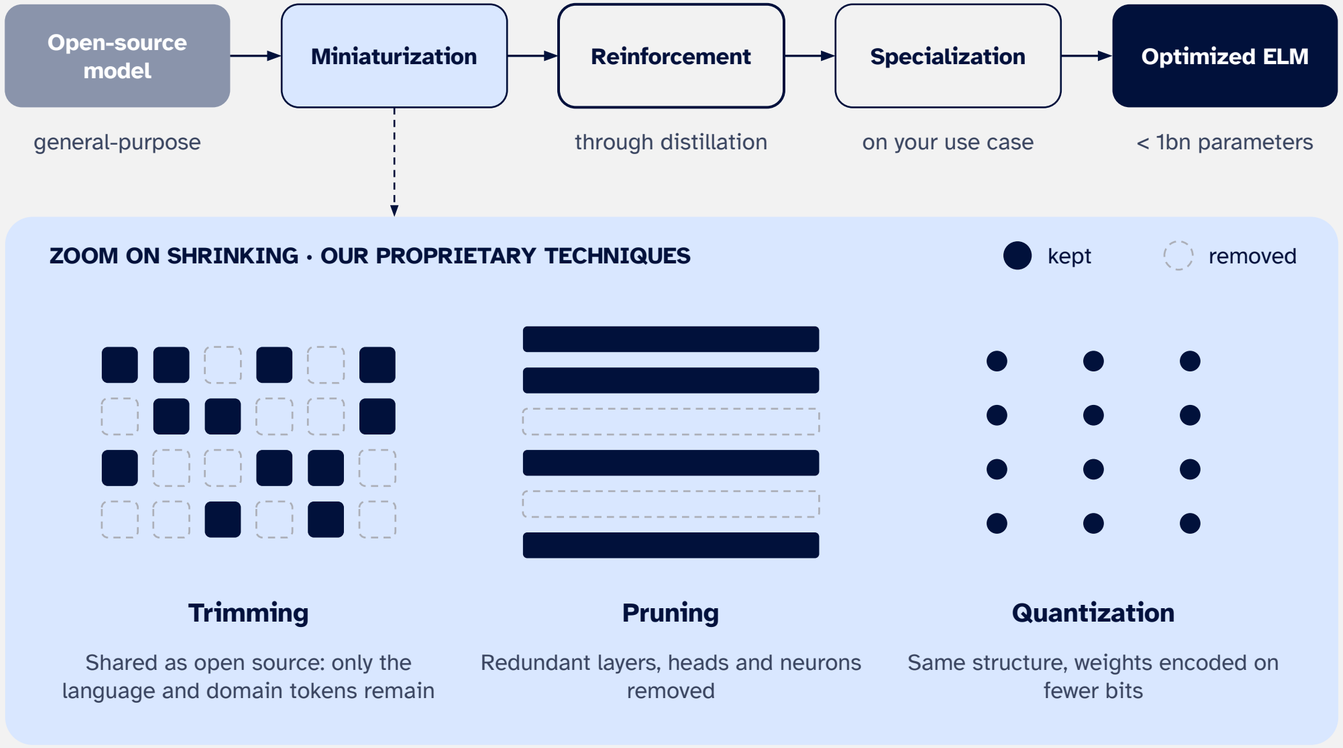
We collect, source and generate the data that makes a model an expert in one or more domains.

RESULT

100% proprietary, < 1bn parameters

Our lab creates optimized AI models

AlphaEdge also optimizes models built on open-source components and specializes them with its proprietary techniques to create **models in its optimized ELM range.**



STEP 1 · SHRINK

Lighten the model

Our proprietary quantization, pruning and trimming techniques shrink the model until it fits in an embedded device.

STEP 2 · REINFORCE AND SPECIALIZE

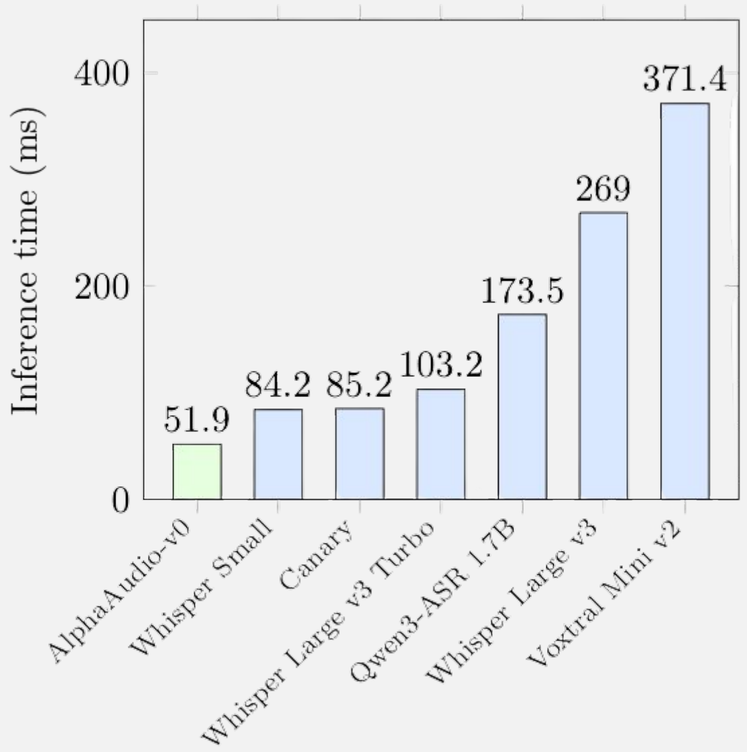
Keep performance, target the use case

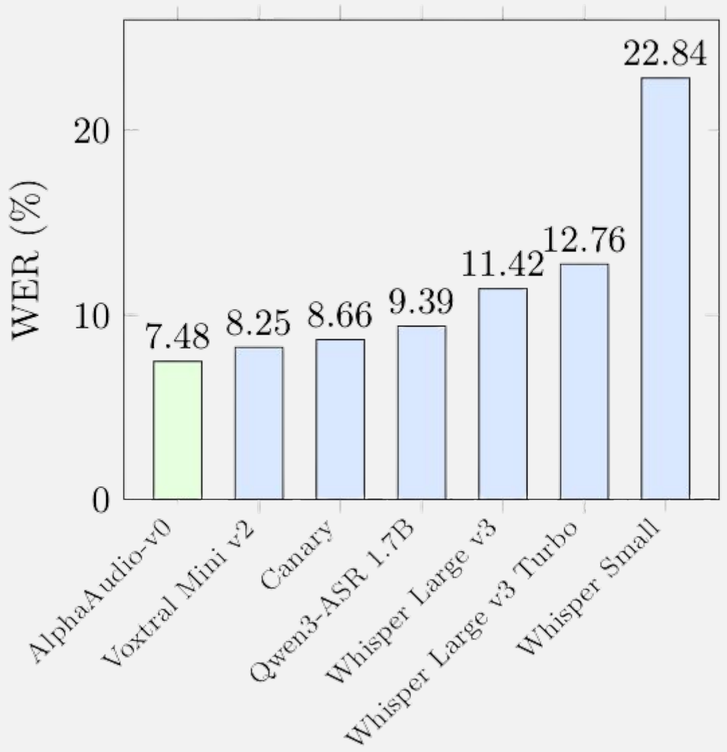
Distillation offsets the performance loss, then our domain data specializes the model for a use case.

RESULT

Under 1 month, < 1bn parameters

Our native audio transcription ELM outperforms the leaders





Full benchmark on alphaedge-ai.com →

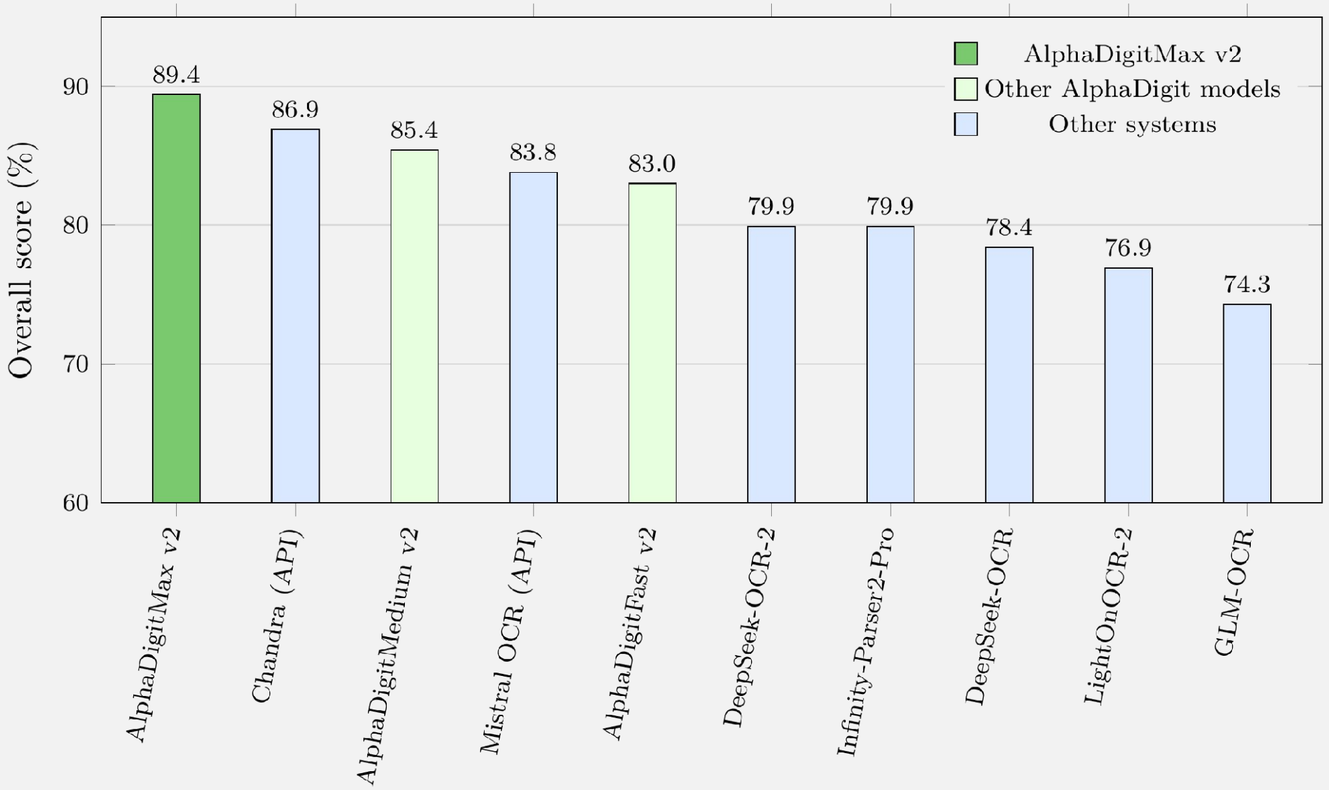
RESULTS

AlphaAudio-v0 evaluated on **Common Voice:**

- **5× faster** than Whisper Large v3
- **WER of 7.48%**, vs 11.42% for Whisper Large v3
- **Global top 10** in French (WER)

Our ASR model is **one of the most efficient in the world for French**.

An OCR model challenging the world leaders



RESULTS

AlphaDigitMax v2 evaluated on **olmOCR-Bench:**

- **No. 1 overall**: 89.4% vs 86.9% for Chandra
- **No. 1 on tables and old scans** (60.6% vs 51.2%)
- **All 3 AlphaDigit models** in the top 5 out of 10 systems

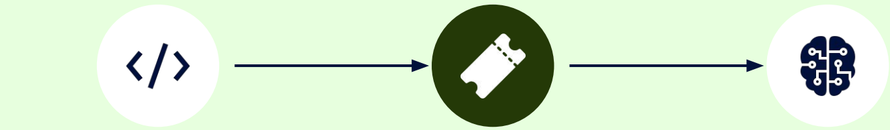
AlphaDigitMax v2 is **the most accurate of the 10 OCR systems compared**.

Full benchmark on alphaedge-ai.com →

Our technology can be consumed **in 2 ways**

01. API consumption

Plug & play model



Developers & software vendors

Toll model: pay per call

Our AI models built into the workflow

PAYMENT MODEL

Pay per use

02. PIPELINES & advanced features

Turnkey solutions by use case



SMEs, mid-caps, large corporates

Turnkey AI chain

Business problem solved

PAYMENT MODEL

One-off fee + monthly license

01. Our AI models available through our API

OCR · VISION



AlphaDigitFast v2

Printed text

OPTIMIZED ELM



AlphaDigitMedium v2

Handwriting

OPTIMIZED ELM



AlphaDigitMax v2

Tables & complex documents

OPTIMIZED ELM

ASR · AUDIO



AlphaAudioFR v0

French transcription

NATIVE ELM



AlphaAudio v1

Multilingual transcription

In design

NATIVE ELM

TEXT · SLM



AlphaLM v0

Text generation

In design

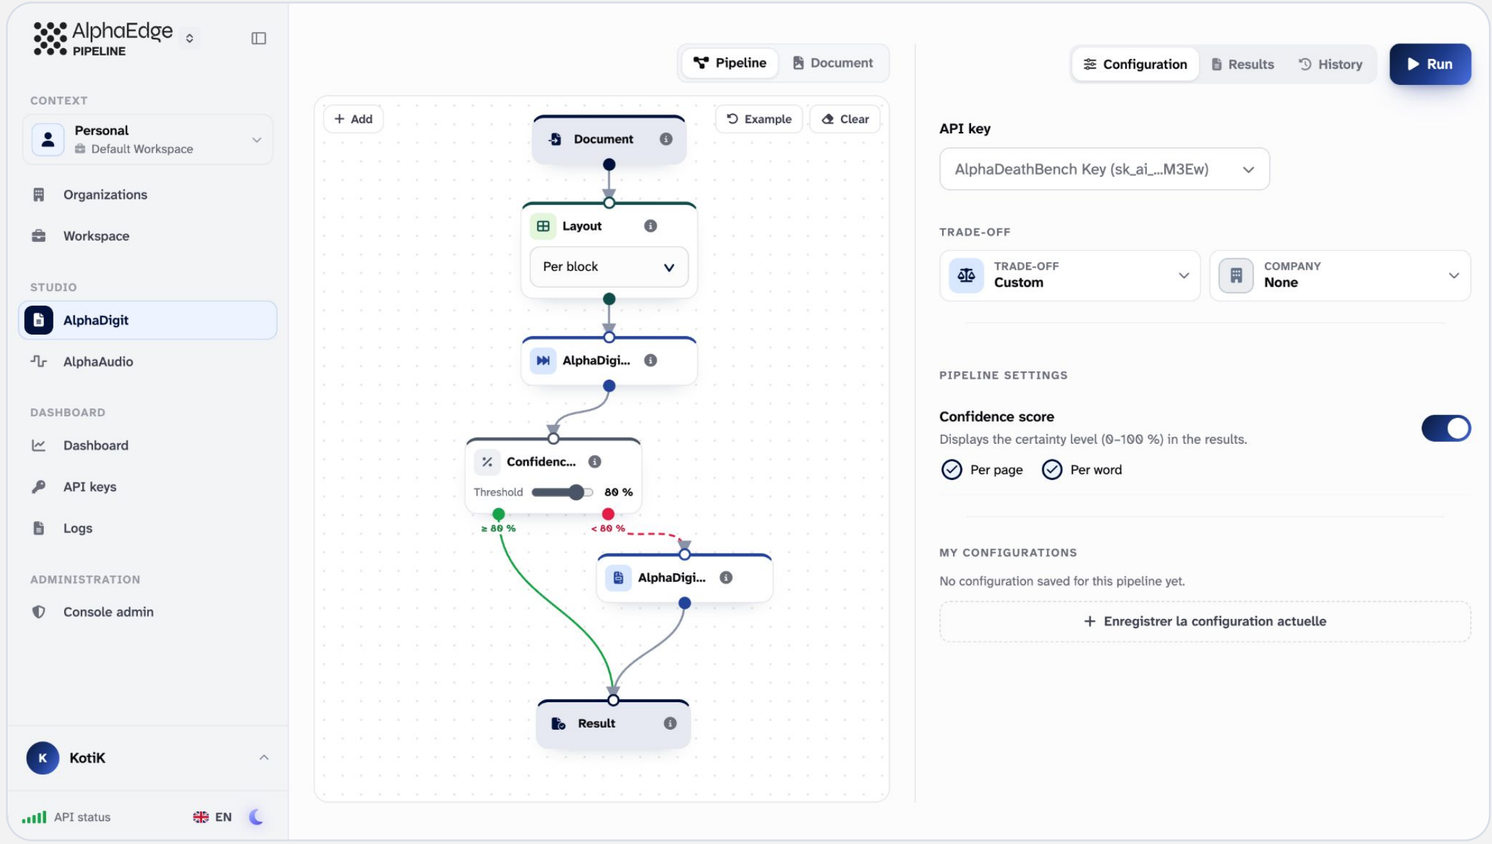
OPTIMIZED ELM

Explore the catalog on our API →

AlphaEdge

02. Our agentic pipelines

Our orchestrator model automatically routes each piece of data to the right AlphaEdge model.



First pipelines deployed

IDP PIPELINE

- **3 OCR models**: handwriting, table, generic
- **3 layout models**: paragraph, line, word
- **1 detection model** for handwriting
- **Confidence scores** and routing rules
- **1 automatic document** sorting model
- **1 targeted information** extraction model

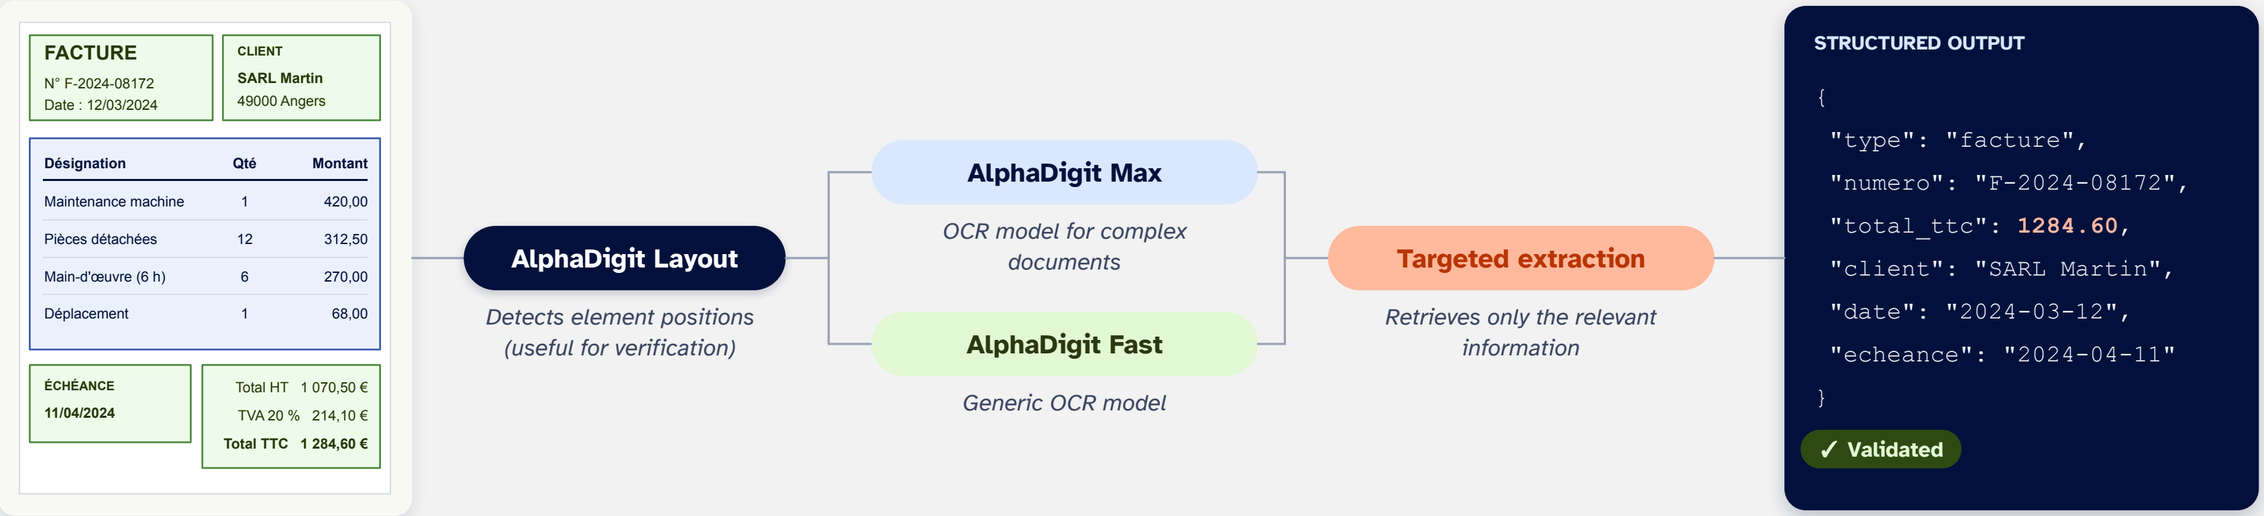
AUDIO PIPELINE

- **1 ASR audio model**
- **1 diarization model:** open source
- **1 post-correction model**

Explore our pipelines →

AlphaEdge

Doxallia case: **98%** of the digitization budget freed up with AlphaDigit (IDP)



OCR QUALITY

World No. 1

on olmOCR-Bench: 89.4%, ahead of the 10 systems compared

SAVINGS

98%

from €15M to €225k per year to process 150M invoices

SOVEREIGNTY

100% sovereign

hosted in France or on-premise: no data leaves

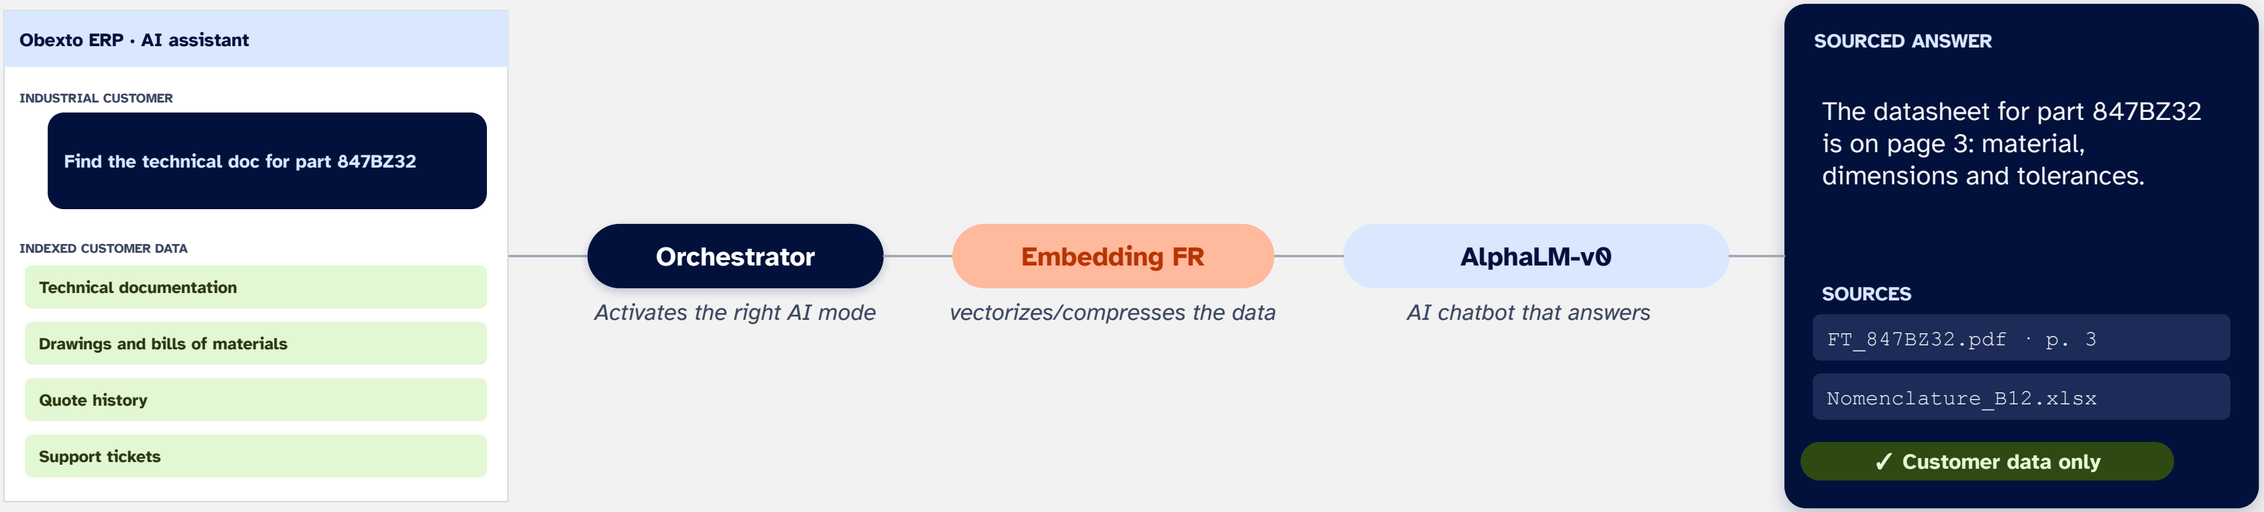
BUSINESS EXTRACTION

Targeted fields

only the useful data, delivered as JSON to Doxallia's IT system

Obexto

Obexto case: **a sovereign AI assistant** built into its ERP software for its 150 industrial customers



QUALITY

No. 1 in French

our embedding model

CONTROLLED COST

÷ 24 vs LLM

cost per token: resale margin preserved

SOVEREIGNTY

Sovereign AI

each customer isolated, data in France

DEPLOYMENT

On-premise

start in the cloud, switch at any time

AlphaEdge

The 4 key steps of a pipeline deployment

01 · 1–2 DAYS

Customer audit, diagnosis and benchmark

From €3k

02 · 2–3 WEEKS

Pipeline customization*

From €5k

03 · 1 MONTH

Deployment and integration

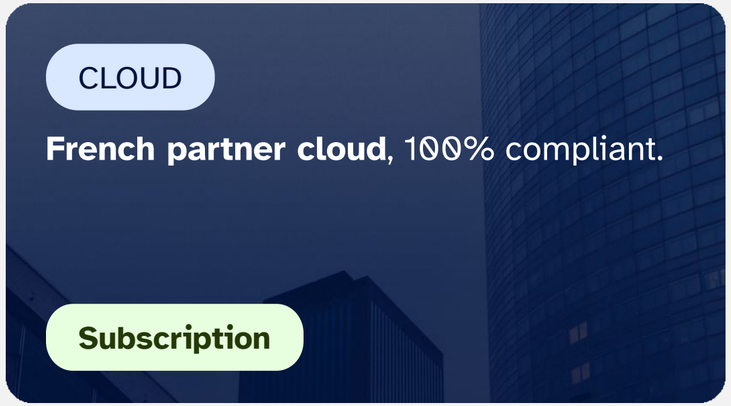
From €5k setup fee

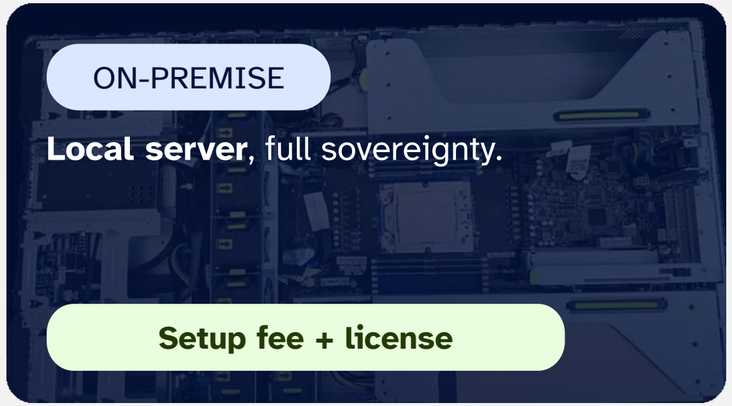
04 · 2–3 YEARS (LTV)

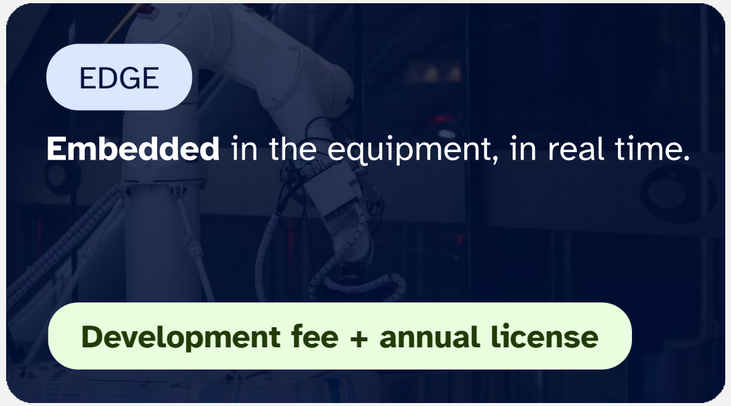
Maintenance, support and model upgrades

From €500/month

3 DEPLOYMENT MODES







AlphaEdge

Commercial traction, **9 months after launch**

SALES STAGE → / PRODUCT LINE ↓	QUALIFICATION	PROPOSAL SENT	NEGOTIATION & CLOSING	POC & SERVICES	DEPLOYMENT
	13 accounts · prob. 10–40%	6 accounts · prob. 50–60%	10 accounts · prob. 70–95%	14 paying customers	3 signed customers
RAG Sovereign knowledge assistants	Cereol, ArcelorMittal, THALES, ONET, TESMEC, WILCO, Les Compagnons du Devoir, Viaduc, PAC LIST	MARINE NATIONALE, NOVAM	Obexto, tgs, NOVAM, GIZEH, SynDev, TALENSIS	Résilians
Embedded AI Embedded AI & vision	CLAAS	POLICE NATIONALE, DATA INCEPTIO, EGLO	wilo, EGLO	Vallemaines
IDP AlphaDigit · document processing	Orange Cyberdefense, leZèbre, zeendoc	Ethypharm, ADENA ENERGIE	DOXALLIA, Luminess	bakertilly, Atos, DOXALLIA	
Audio & Forecast			EGLO	Wattflow, Publicis	
INDIRECT CHANNEL · RESELLERS **Integrators**	sopra steria, Square management	clever cloud	arkan		dri

€195k signed

€1.07M
Signable pipeline: 17 open proposals

€415k
2027 revenue already backed by these deals

€1.27M
36-month revenue potential, open and signed

AlphaEdge

We aim to become the European leader in frugal AI

3 pillars: **efficiency, domain specialization, self-orchestration**.

	AlphaEdge	Mistral AI	MULTIVERSE COMPUTING	Liquid	sapient intelligence	webAI	CALLOSUM
European sovereignty	✓	✓	✓	–	–	–	✓
Efficiency & < 1bn parameters	✓	–	–	✓	✓	✓	–
Self-orchestration	✓	–	–	–	–	–	✓
Domain specialization	✓	–	–	–	–	✓	–
General knowledge	–	✓	✓	–	–	–	✓

An HR structure built for international expansion

21 → 143 FTEs

from end-2027 to end-2030, 49% in R&D and engineering

2027 · PRE-SEED

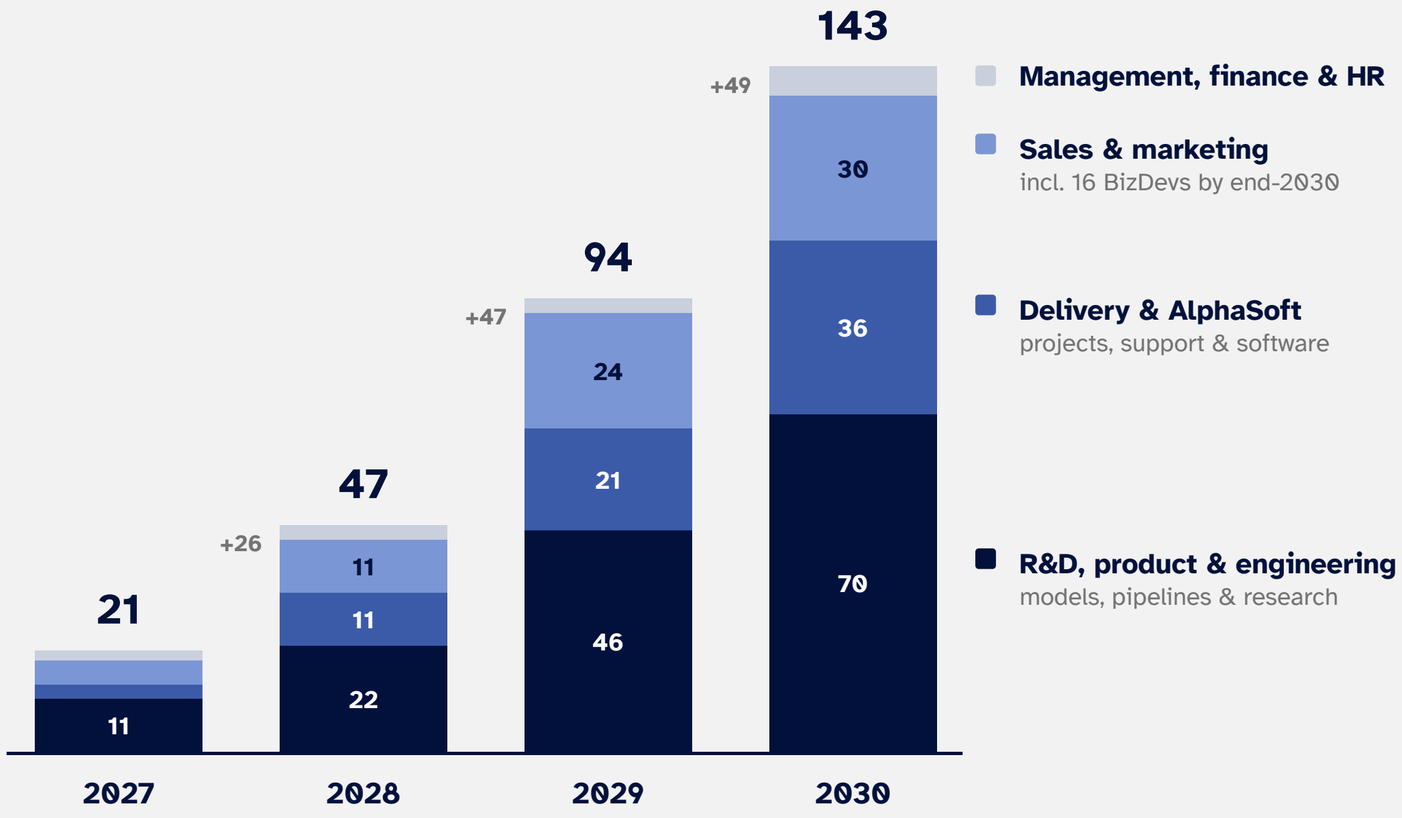
Build the sales and delivery teams

Senior ML engineer in April, 2 senior BizDevs in July, a 3-person delivery team.

2028 → 2030 · SCALING UP

Structure the growth

BizDev hiring ramps up from June 2028, AlphaSoft BU in 2029. Over €300k revenue per FTE targeted for 2030.



Year-end FTEs · Business plan, base case

AlphaEdge

Industrializing customer acquisition

3 levers and a short cycle that secures customer acquisition through a paid audit or POC.

DIRECT

Automated hybrid outbound

Marketing, sales and agentic AI synergy

≈ €1,250 per qualified lead

INDIRECT

Partners & integrators

White-label IT services firms and integrators

≈ 15% commission, success-based

MARKETING

Ecosystem & inbound

Trade shows, press and networks

€280 per inbound request

FULL SALES CYCLE · 3 TO 6 MONTHS

Qualification & demo → **Paid audit or POC** → **Pipeline project & recurring**

≈ €23k
Average CAC in 2027

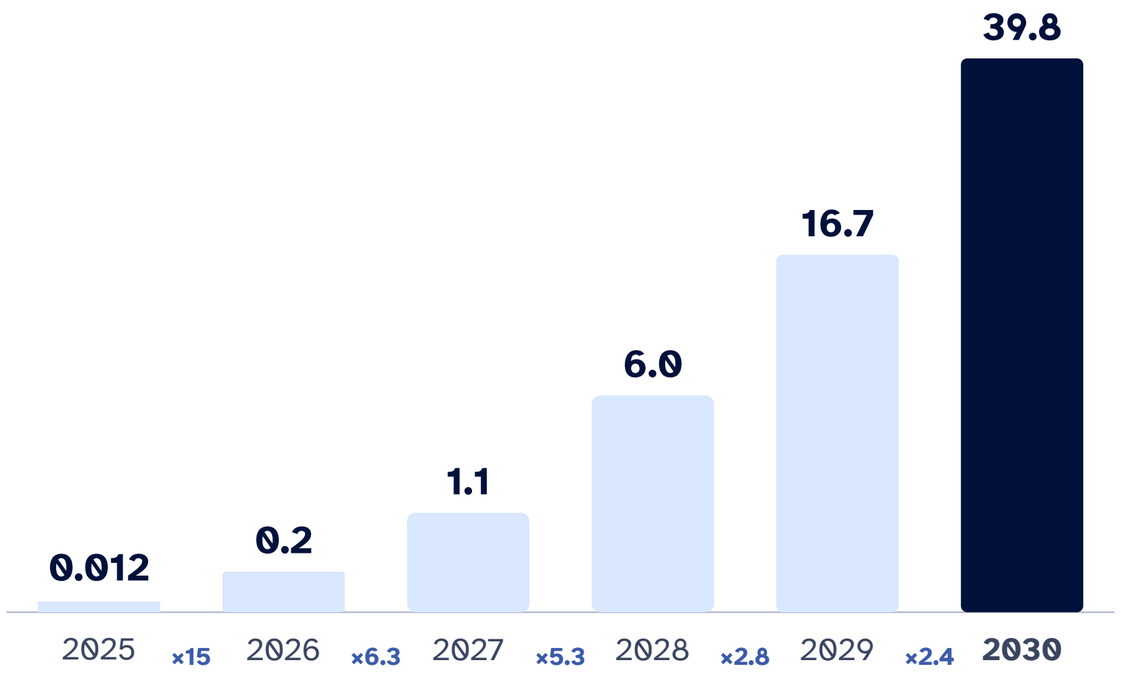
4.5
LTV / CAC in 2027 (36-month contract)

9 months
CAC payback

90%
of customers start with an audit or a POC

Growth: **€39.8M revenue in 2030**

Revenue in €M · BP of 15 Nov. 2026, base case · compressed scale for readability



€19.0M

ARR by end-2030

91%

gross margin

2028

first positive EBITDA

17%

of 2030 revenue from international

AlphaEdge

Technical excellence serving the business



Théo Hubert

CTO & CO-FOUNDER

AI architecture · data · strategy

200
engineers managed

7 years
of entrepreneurship

Dozens
of validated MVPs

AI engineer & researcher · ESAIP & Stanford



Fabien Fasson

CEO & CO-FOUNDER

Sales · go-to-market · strategy · fundraising

+€6M
B2B SaaS revenue

7 years
selling to large corporates

3
projects to product-market fit

Business management · Paris 1 Panthéon-Sorbonne

Our organization, **backed by our mentors**

MANAGEMENT



Fabien Fasson
CEO & co-founder
Sales, strategy, fundraising



Théo Hubert
CTO & co-founder
AI architecture, data, engineering

SENIOR MANAGERS



Célia
CMO



Alexandre
R&D Director, PhD



Loïck
Lead Data Scientist

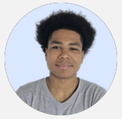


Serge
PhD, deep learning

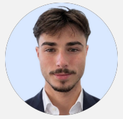


Marie
PhD, mathematics

TEAM



Dariel
Business developer



Esteban
Business developer



Mattéo
AI engineer



Léo
AI engineer



Léo
Full-stack dev



Théo
Full-stack dev

MENTORS

FC — **Frédéric Clabaud**
Angers Technopole

JB — **Jean-Philippe Bellaiche**
AxelNova

SM — **Sylvain Maussion**
Emsylado

VB — **Violaine Berté**
Olbee

An ecosystem that funds us and opens doors

TECH PARTNERS

LERIA

Le Mans Université

OVHcloud

Scaleway

EASIO-COM

NVIDIA INCEPTION PROGRAM

LABEL

LABELISED DEEP TECH

La FRENCH TECH bpifrance

PARTNERS & INSTITUTIONS

RÉGION PAYS DE LA LOIRE

CCI MAINE-ET-LOIRE

esaip École d'Ingénieurs La Salle

L'incubateur ANGERS TECHNOPOLE

ATLANPOLE Land of innovation

FRANCE 2030

clever cloud

CREDIT COOPERATIF

incubateur BELLE DE MAI

Images & Réseaux

ID4 MOBILITY

Réseau Entreprendre

Réseau Initiative VAR

What AlphaEdge has **already built**

INTELLECTUAL PROPERTY

1 patent

granted, 2 more in preparation

FUNDING

+€900k

already secured, including **€250k from the founders**

TALENT

13 people

including researchers and data scientists of world renown

DATA

+5 M

documents, one of the largest multilingual handwriting datasets

INFRASTRUCTURE

+€150k

of hardware installed in a Paris data center

VISIBILITY

+1 M

downloads on Hugging Face, 2 TV appearances and national press

A **€500k** pre-seed for efficient AI

€350k already raised through BSA-AIR convertible notes. Now raising €500k in equity for January 2027.

EQUITY

€500k

equity round, January 2027

LEVERAGE

€0.7M

Bank loans, Bpifrance, convertible bonds

R&D & TECHNOLOGY

€59k

- Senior ML engineer in April
- GPUs, servers & data
- 2027 R&D budget: €1.42M

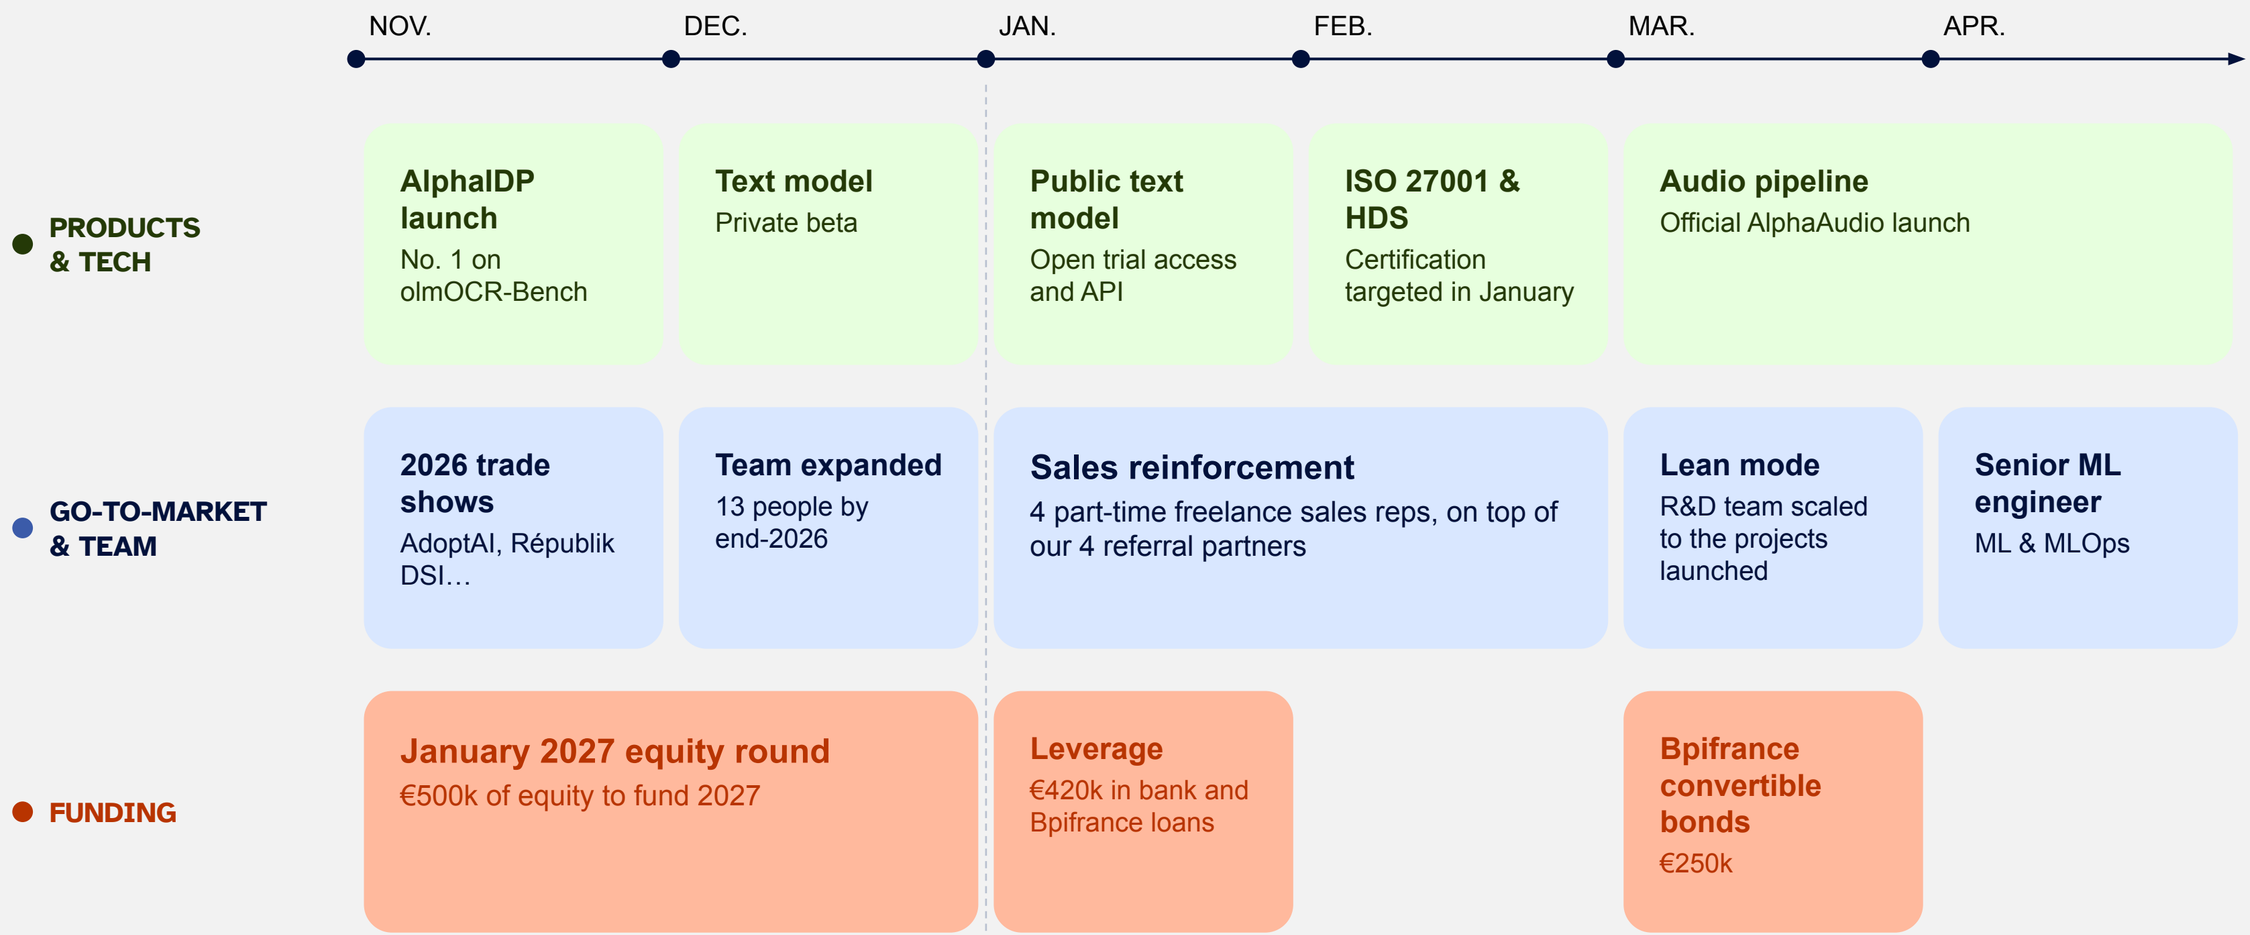
COMMERCIAL EXPANSION

€441k

- 2 senior BizDevs in July
- Freelancers & referral partners
- Marketing & 6 trade shows

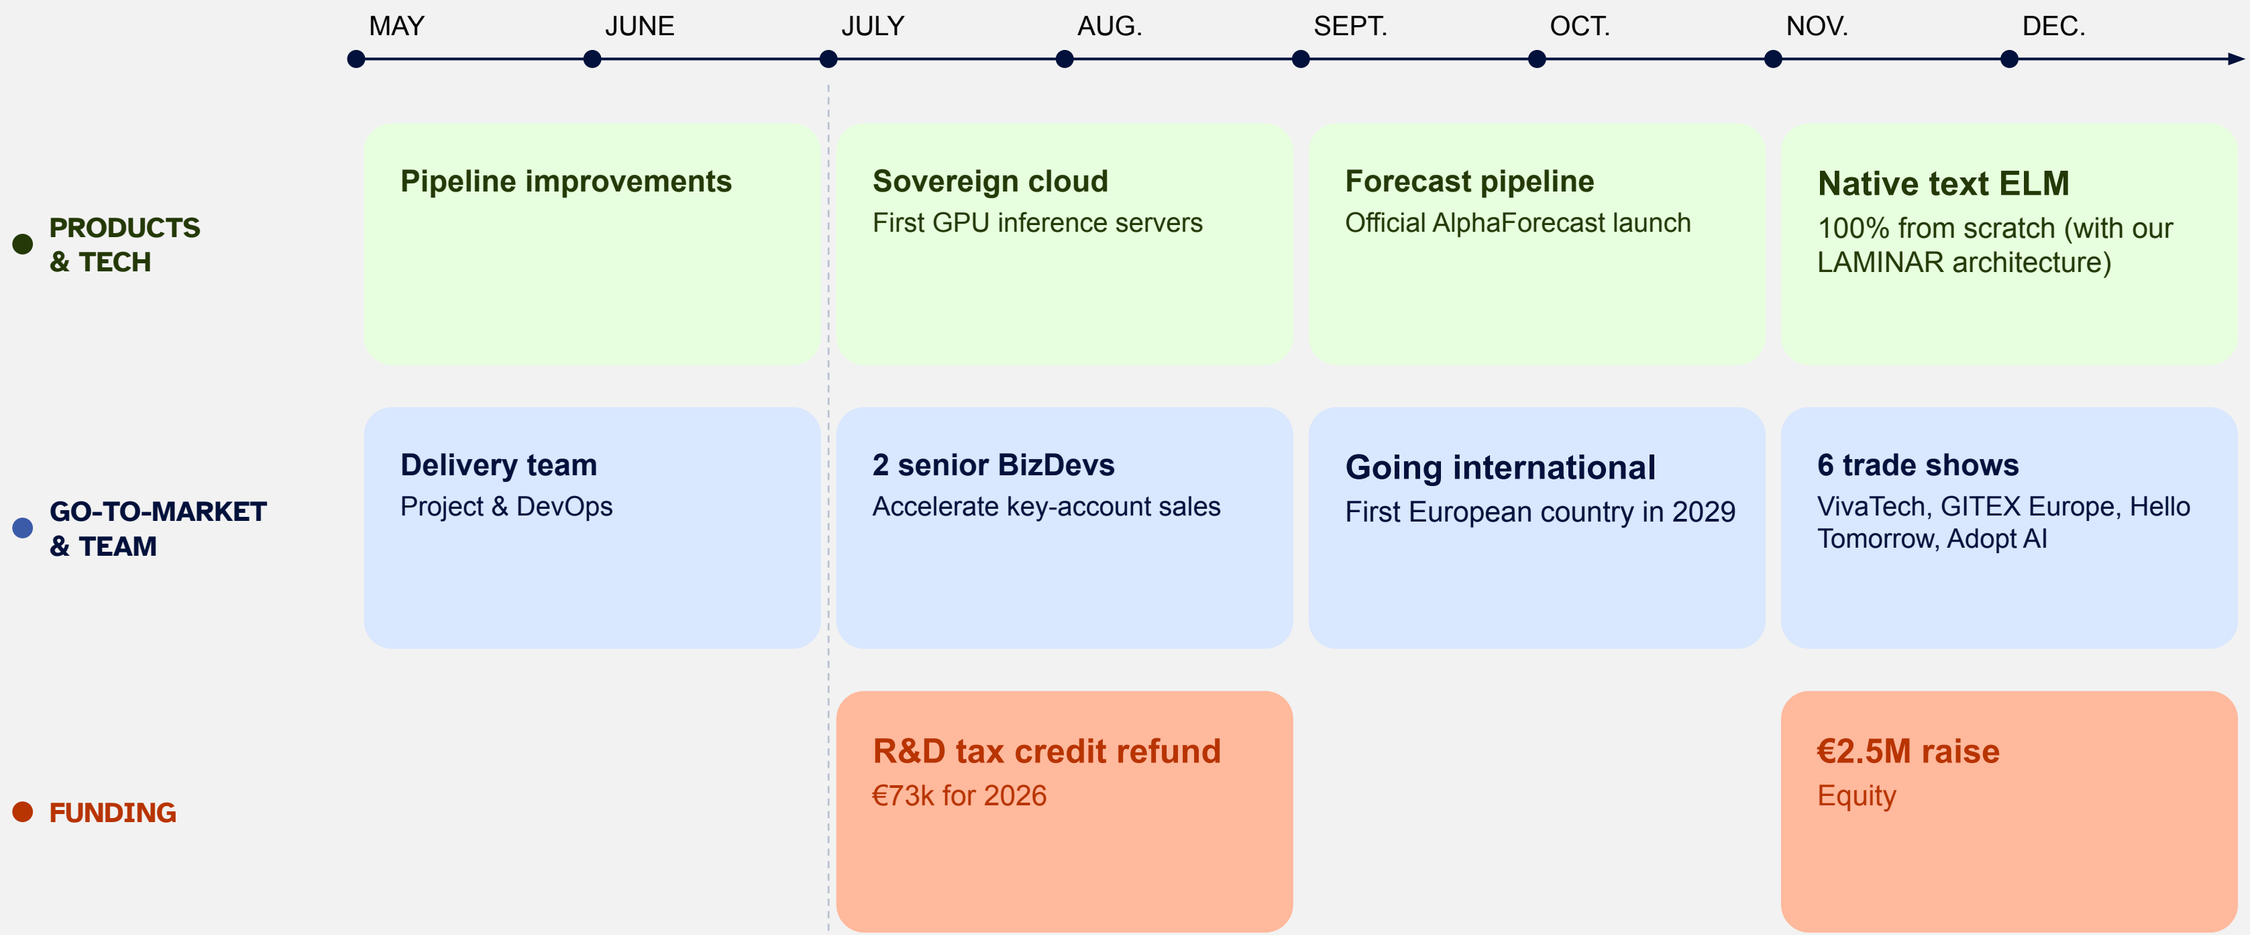
MILESTONES 15 customers by end-2026 · €1.1M revenue, €274k ARR, 21 FTEs by end-2027 · **€2.5M raise in Dec. 2027**

Roadmap: **November 2026 → April 2027**



AlphaEdge

Roadmap: **May → December 2027**



AI technologies with ESG impact

Cut our users' GHG emissions by **70%**

Cut water consumption by **70%**

10× less IT hardware needed

Make AI **affordable** for public services

Foster **job creation** across the country

Achieve **gender parity** in leadership roles and **share value** internally (employee stock options)

AlphaEdge

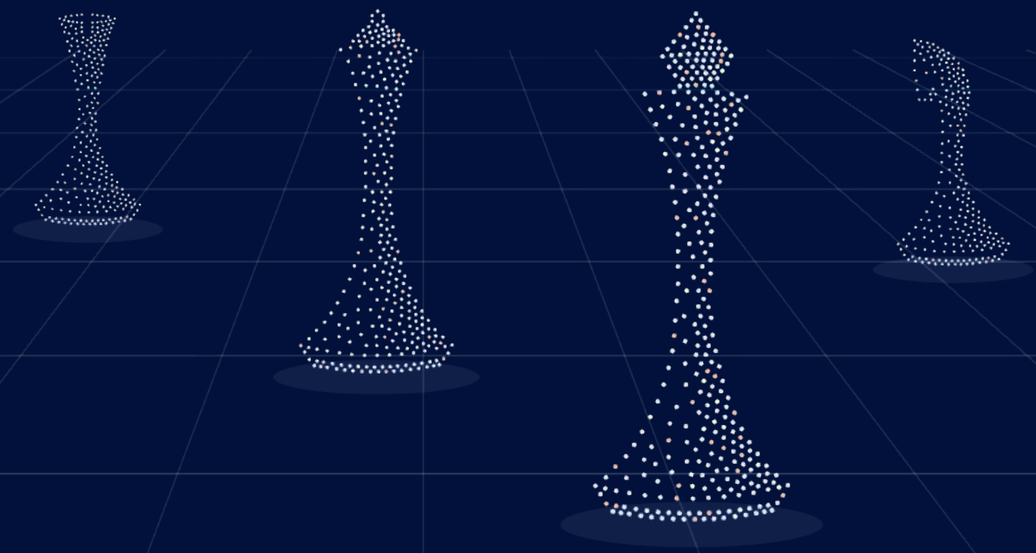
Building AI where it matters

CONTACT

Fabien Fasson, CEO

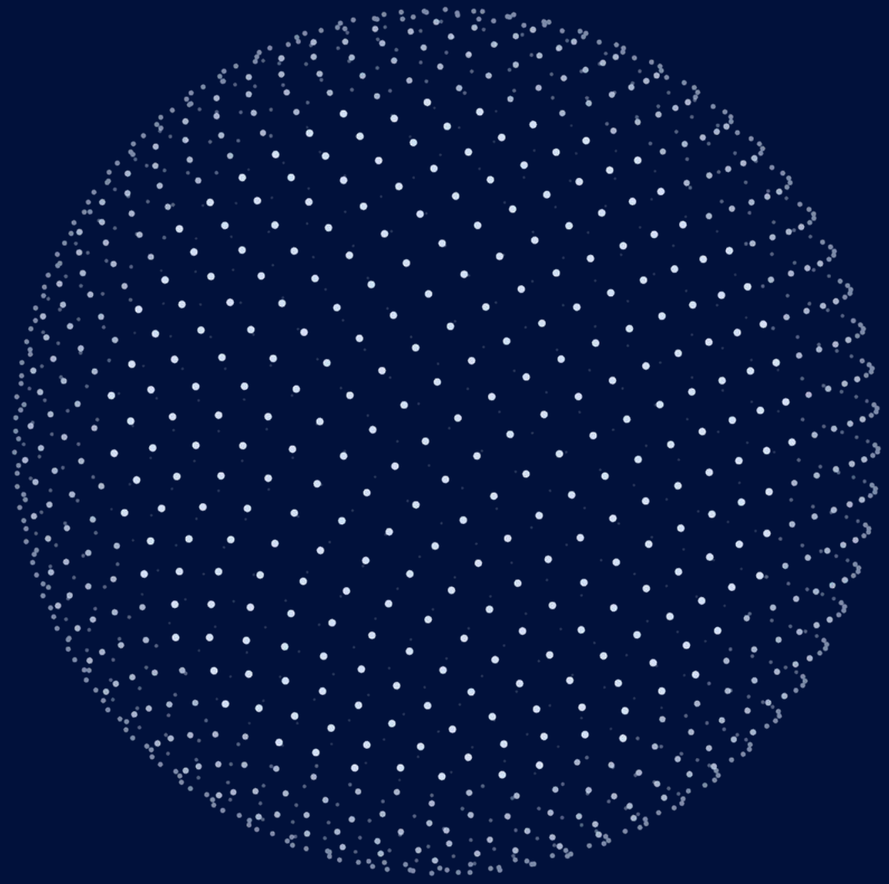
+33 6 88 69 53 00

fabien.fasson@alphaedge-ai.com



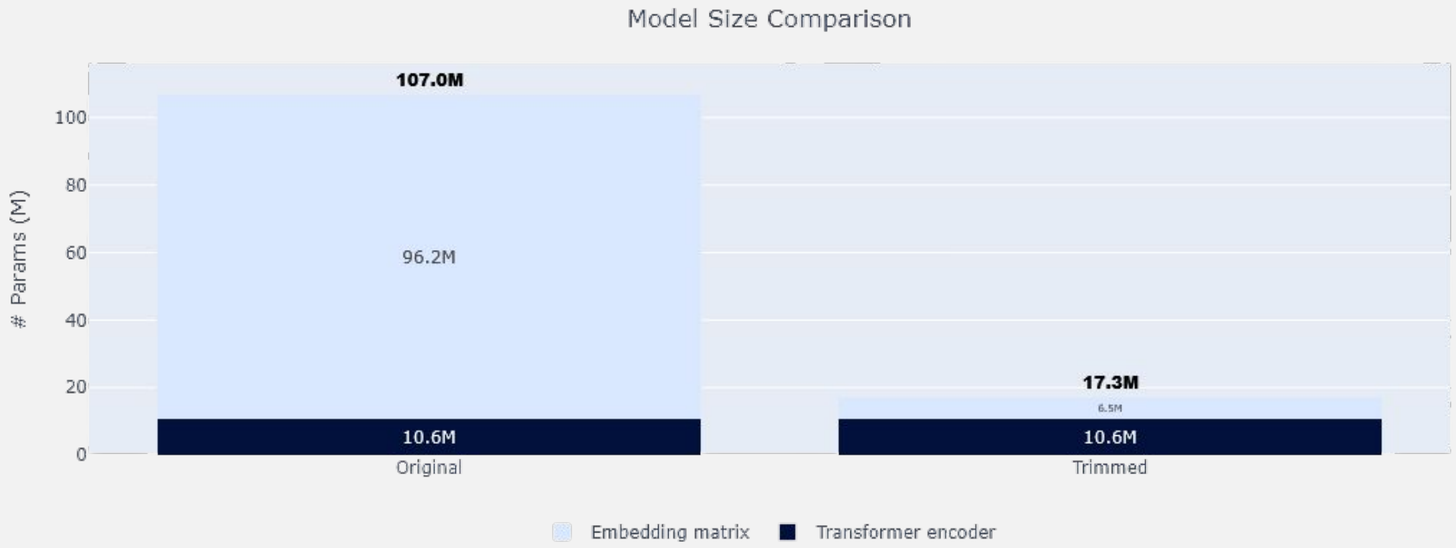
Appendix

PIPELINES · COMPRESSION, DATA, INFRASTRUCTURE · ALPHAPEN · FINANCIALS



Cutting-edge expertise in compression

Our ELMs address digitization and transcription needs. To broaden our scope, **we complement our expertise with the compression of existing models**, notably through **trimming**. This way, we do not depend solely on our own architectures and we strengthen the pillars of our value proposition: **efficiency, specialization and sovereignty**.



Cutting-edge technology

Distillation, quantization, pruning, trimming: we master a broad range of state-of-the-art techniques.

Open-source recognition

About 5,000 open-source models optimized through *trimming*, a little-known and highly effective technique.

Example: *ibm-granite/granite-embedding-107m-multilingual*, optimized for Arabic.

End-to-end data mastery

01

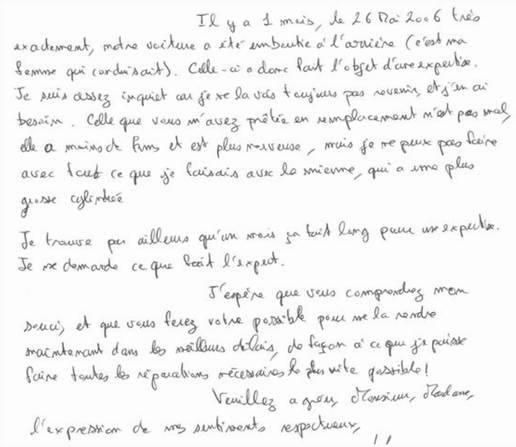
Field data collection



Handwritten forms collected in the field, anonymized and used only to train our AI.

02

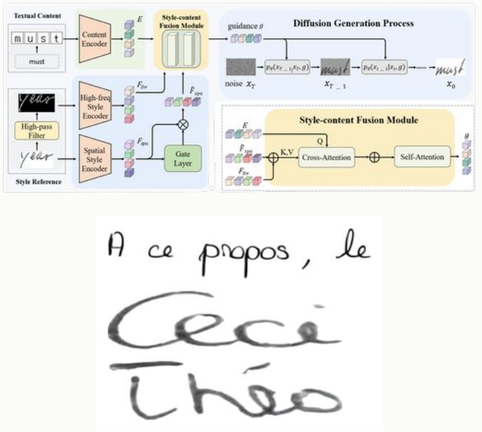
Sourcing of openly licensed data



Loïck Bourdois, Hugging Face Fellow, is internationally recognized for his data expertise.

03

Synthetic data generation



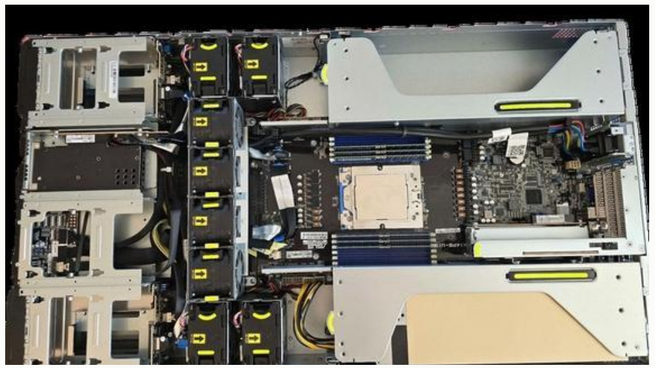
We trained our own models to generate part of our synthetic data.

AlphaEdge

From training to inference, all in-house

04

Training on our own server

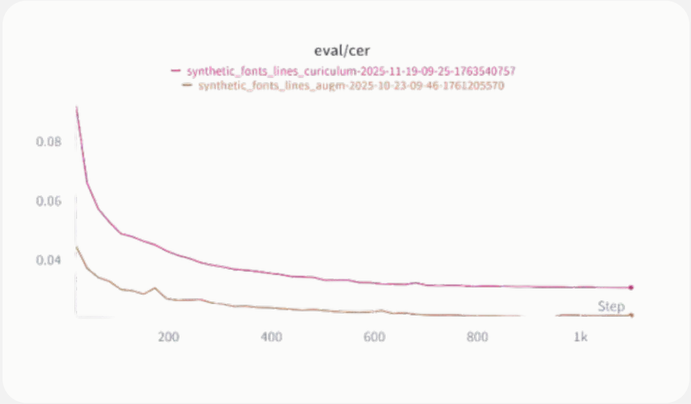


4× RTX PRO 6000 · 400 GB RAM · 6 TB · AMD EPYC 32 cores.

Scaleway startup program: €3k of cloud credits per month.

05

Model validation



Evaluation curves tracked at every training iteration.

06

Production deployment on our inference server

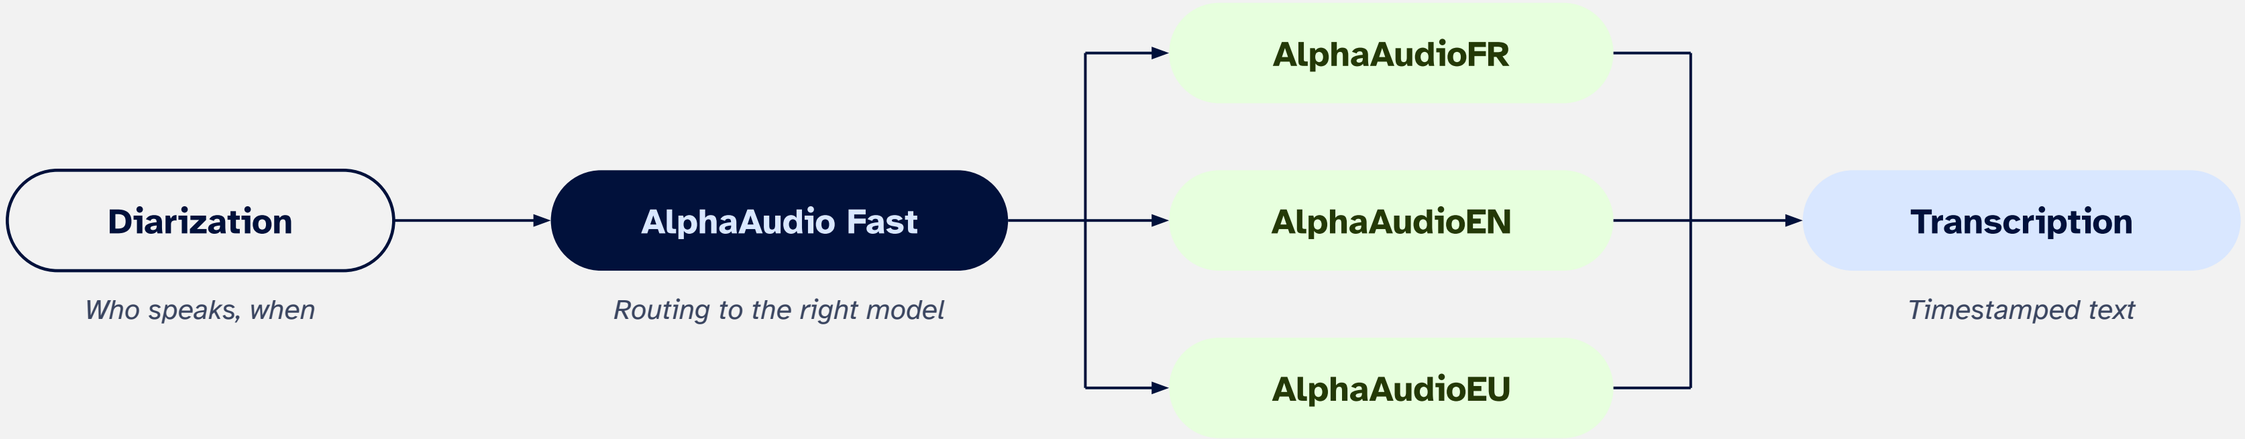


1× RTX 5090 32 GB VRAM · 32 GB RAM · 1 TB.

Serverless overflow (Modal, EU region, GDPR-compliant) for usage peaks.

Adaptive audio pipeline

The right audio model for each use case. Faster, cheaper, **with no compromise on performance**.



USE CASES

Transcribed meetings
Meeting minutes

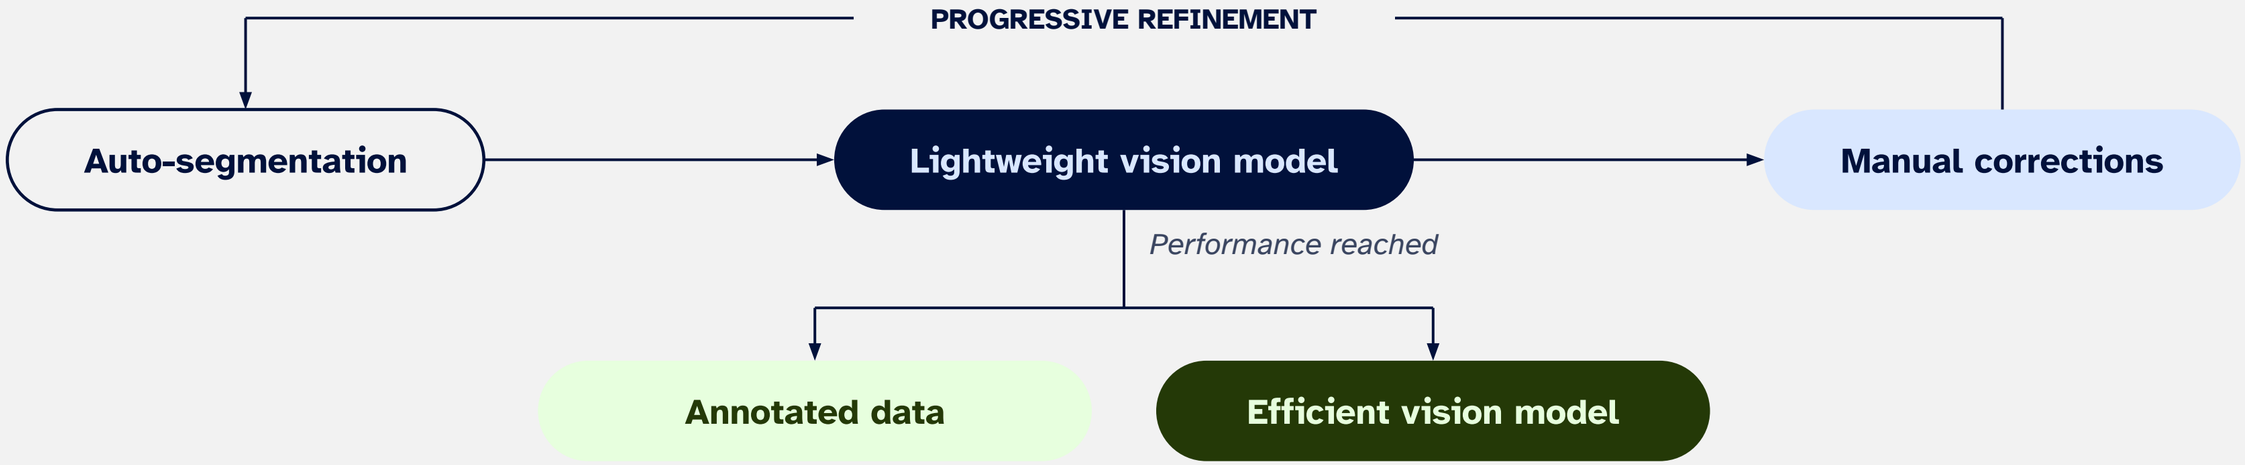
Customer relations & after-sales
Call analysis

Indexed content
Podcasts, videos, archives

Voice control
Embedded, offline

Iterative vision pipeline

Frugal vision, trained on your data until it reaches the target performance.



USE CASES

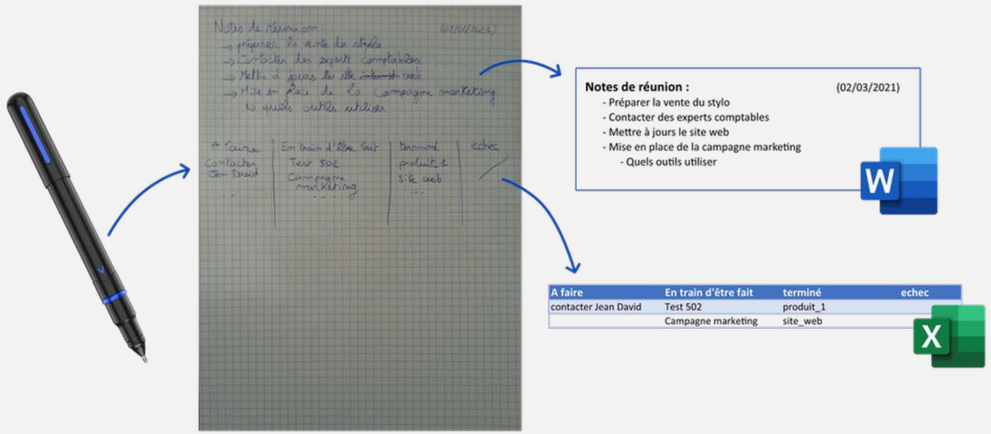
Quality control
In-line micro-defects

Predictive maintenance
Wear and anomalies

Precision agriculture
Harvest and monitoring

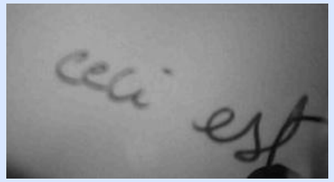
Traceability & security
Flows, inventory, sites

Our first use case was born from embedded AI



We designed a **pen that digitizes handwriting in real time**, on any surface.

AI WORKFLOW · ALPHAPEN



AI < 60M · CER < 4%

Embedded digitization

AI < 50M

Embedded OCR post-correction

ALPHAPEN, A STRATEGIC ASSET

- **1 granted patent, 1 registered trademark**, 1 embedded AI chain dedicated to handwriting
- **2,000 pre-orders** (€300k crowdfunding potential)
- A pre-industrial prototype built with **€150k of investment** and **6 years of R&D**

Projected income statement **2026 → 2030**

Positive EBITDA from 2028, €9.9M net income in 2030.

€k excl. VAT	2026	2027	2028	2029	2030
Revenue	**182**	**1,146**	**6,045**	**16,683**	**39,835**
Growth	n.m.	+530%	+427%	+176%	+139%
of which recurring revenue	7	159	2,237	8,284	22,103
Gross margin	94.4%	89.1%	89.6%	89.2%	90.7%
Personnel costs	351	883	3,294	7,536	11,671
EBITDA	**–282**	**–343**	**509**	**3,611**	**14,204**
EBITDA margin	–154.9%	–29.9%	8.4%	21.6%	35.7%
R&D tax credits	73	211	511	1,394	2,868
Net income	**–212**	**–304**	**685**	**3,112**	**9,934**

Source: AlphaEdge business plan, updated 15 November 2026, base case.

Projected cash position **2027 → 2030**

Positive cash throughout the plan: €40.0M by end-2030.

€k excl. VAT	2027	2028	2029	2030
Cash flow before financing	**–1,172**	**–2,885**	**–4,956**	**6,553**
of which capex	–633	–2,903	–8,455	–6,951
Average monthly burn	98	240	413	–
Equity raised	3,000	9,500	–	–
Financial debt at year-end	1,299	7,274	23,030	31,261
Lowest cash point	403	4,232	20,447	30,859
Months of runway at low point	5.3	10.1	15.6	11.8
Months of runway in December	15.3	24.2	15.6	13.2
Year-end cash	**3,052**	**15,584**	**26,005**	**40,021**

Source: AlphaEdge business plan, updated 15 November 2026, base case.