




Every conversation a company has should be intelligent, instant, and always available - without losing the human touch where it matters.

Yara - Autonomous Voice Infrastructure

Nobody should have to choose between scale and quality.

Today they do. Either you hire endlessly and still miss calls, or you automate with rigid systems that frustrate everyone who touches them.


We exist to end that trade-off - to make every conversation autonomous and excellent.



The world changed - conversations didn't keep up.

- Customer expectations moved to instant, 24/7, no queues.
- Call volumes are exploding. Teams and budgets are not.
- Voice AI crossed the threshold - good enough to trust in production.

The gap between what customers expect and what organisations can deliver has never been wider. That gap is the opportunity.



Hiring more people scales the cost, not the solution.

01

Repetitive calls

A large share of every team's time goes to calls that never needed a human.

02

Unanswered calls

Nights, weekends and peaks go unanswered - lost revenue and lost trust.

03

Lost context

Every handoff loses context. The customer repeats themselves and quality drifts.

04

Headcount

Adding people multiplies the cost of the problem without removing it.

Insert your own figures here before sending - this is where the deck lands or doesn't.

The gap

What a demo shows you. What production asks for.

ABOVE THE LINE - WHAT YOU SEE

A prompt
A demo call
A voice on a workflow tool
A handful of scripted paths

BELOW THE LINE - WHAT IT TAKES

- Real-time voice orchestration
- Barge-in and interruption handling
- Telephony edge cases and transfers
- Identity verification and one-time codes
- Multi-agent routing on shared context
- Guardrails enforced server-side
- Fallbacks, retries and timeouts
- Audit trail and data residency
- Evaluation at scale before go-live

The demo is the easy tenth.



One vendor, one price list, and no way out.

01

A price you don't set

Per-minute and per-token pricing is decided by someone else. Welded to a single provider, you have nowhere to move when it changes.

02

Today's best model isn't next quarter's

Latency, comprehension and naturalness shift every few months. A stack tied to one vendor runs last quarter's quality until someone rewrites it.

03

Demos are built on the easy cases

Restaurant bookings and property viewings are not regulated. Identity, money and medical records are. Proven on the easy path is unproven on yours.

So we orchestrate instead of reselling - and we pick per use case, not per vendor.

The pivot

Two ways to run your conversations.

THE OLD WAY

Humans handle every call

Limited to business hours

Cost grows with volume

Context lost at every handoff

Rigid phone menus

WITH YARA

AI handles the routine, humans handle what matters

24/7, instant, no queue

Scales without scaling headcount

Context retained across the whole call

Natural conversation, in your customer's language

How we're different

A platform, deployed like a partner.

Voice Orchestration Platform

We orchestrate the best voice models on the market instead of betting on one. You always run the state of the art, and you are never locked into a single technology.

Forward Deployment Company

We don't hand you a tool and leave. Our engineers embed inside your organisation to integrate, calibrate and run the system until it works - as an extension of your team.



The power of a platform. The certainty of a partner.

We don't bet on one model.
We benchmark them all.

Continuous benchmarking

We test the best voice models on the market against precise quality indexes: latency, naturalness, comprehension, reliability.

The right model per job

We orchestrate the right model for each use case and swap as the field moves. You run the state of the art without ever migrating.

Matched to your needs


We align each model's strengths with your specific job - never a single default applied to everyone.

The field moves every quarter. You shouldn't have to.

What we measure, before you hear it.

Every model we consider is scored on the same six axes - and scored on your own recordings, not on a public leaderboard.

- Scored on your calls, not a leaderboard.
- Re-run each quarter, as the field moves.
- A winner per use case, not per vendor.



It doesn't just talk. It understands.




Knowledge graph and embeddings

The agent knows your business, your products and your rules - not just keywords. It answers correctly because it actually knows.




Context accumulator

It captures and uses the full context of the conversation in real time. Your customer never repeats themselves - like your best human agent, at scale.



One conversation. A team of agents behind it.

Specialised agents collaborate through a shared context, so even a complex call is handled end to end.



THE SPECIALISTS BEHIND IT

- Identification and one-time codes
- Policy and customer data
- Ticketing and routing
- Escalation to a human
- Compliance and guardrails
- One shared context across all of them

Autonomous - but never unaccountable.



Human in the loop

For any critical operation, a notification lands on a human's desk to approve or reject before anything happens.

Approve

Reject



Security and authentication

Identity verification and one-time codes handled live during the call. Auditable, and enforced server-side - never in the prompt.



You stay in control of what matters. The AI handles the rest.

We don't hope it works. We prove it first.



300+

scenarios before go-live

Edge cases, difficult callers, adversarial inputs.
Every agent is stress-tested before it takes a
real call.



AI testing AI

We simulate thousands of calls with AI-driven
callers against your agent, at a scale no human
QA team could reach.



Joint calibration

A phase run with your team, tuning the agent
to your reality. It is the difference between a
demo and production.

No blind spots. No surprises. No breach in the agent's behaviour.

Confirm the scenario count against your own QA records before sending.

Proof

This isn't a vision. It's in production.

LIVE

First autonomous AI call center running at a regulated insurer in the GCC.

CLEARED

Passed a health authority's review. Scoped so sensitive data stays out of the AI path, with human handoff on anything critical.

REAL


Carrier-grade telephony, in-country data residency, genuine concurrent call volume.

If it holds under a regulator, it holds for you.

What we build on

We build on the best.
We're owned by none of it.

 **CARTESIA**

 **Fish Audio**


 **ElevenLabs**

 **Qwen**



Speechmatics

 **Neuphonic**

Our value isn't a single vendor. It's the layer that picks
and combines the best for you.



From your systems to autonomous, in weeks.



The first step

Start small. Prove it. Then scale.

Scoped

One department or one use case. Defined scope, real conditions.

One metric

Agreed up front: the share of conversations handled autonomously, with zero critical escalation missed.

Reversible

It sits in front of your setup, not inside it. Switch it off at any time.

Low risk by construction. You can undo it in an afternoon.

The team

The people behind Yara.



Moussa Touré

Founder & CEO
Ex Microsoft / +10y of experience



Soma Bini

CTO. Previously engineering
manager for Autonomous Cars



Madriss Seksaoui

Chief AI Lead +10y of XP.
Deployed Gemini for 300 000 users



Mouctar Diallo

Senior Full Stack



Richard Bathiebo

Senior Full Stack



Sabine Di Mattia

Chief Operating Officer



BACK TO WHY WE STARTED

Every conversation your organisation has can be intelligent, instant and always available - without losing the human where it counts.

That's not the future. It's ready now.

Let's run your pilot. moussa@yaracx.com · +971 58 557 9423