LEAP

RIYADH, 2026

ISSUE 001

THE LEAP EDITION

sawt.sa

The Customer Interaction Company

SAWT



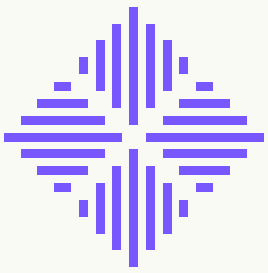


Table Of Contents

MANIFESTO
Craftsmanship at the Frontier

ACT I

THE PROBLEM
How Complex a Phone Call Is
The Whole Surface
The Real Cost of a Call
CX Standards in the Region

ACT II

THE BUILD
Understanding Context
The Four Suites
Buy or in-house
The Best Arabic Models
The Signal Loop
Scalability

ACT III

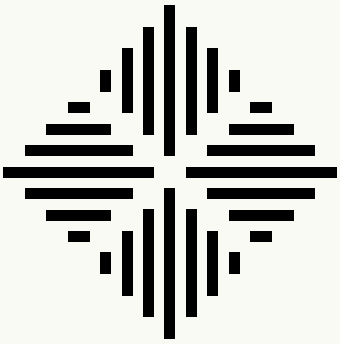
THE PROOF
The Numbers
The Work
Compliance

ACT IV

INSIDE SAWT
Praxis
The Team, and the Obsession
The Eight Values
Real AI

Craftsman

At th



Company

nship
e Frontier

Every business is remembered through its customer interactions. Too many of those memories are bad ones: the call that never connects, the story repeated to a third department, the complaint that disappears.

The people handling them care. The operation around them is simply one of the hardest in business to run well, and when demand outruns capacity, quality gets rationed. Somewhere along the way, poor satisfaction became a statistic businesses learned to live with. That acceptance is our opposition.

Everyone here knows what a good interaction feels like. The pharmacist who knows your name. The voice note answered in minutes, in your own dialect. Personal, immediate, resolved. That standard belongs to every customer of every business, at any scale. So far, only businesses small enough to know their customers by name have managed to keep it. For the first time, closing that gap is possible,
and our mission is to turn these operations into software.

The foundational layer is open: intelligence anyone in the world can build on. We build software that is personal and contextual. It answers in your dialect, knows your history, understands why you reached out, and resolves issues on the first try.

Context that is solved through the expectations of this region and its customers. It is learned through production, refined at scale, and improved with every interaction. That flywheel keeps us at the frontier. It is why the best Arabic models in the world are ours. We hold ourselves to the standard we are building toward. We sell outcomes, aggregated with taste: which is the care taken over every detail a customer touches. Sophisticated frontier technology comes from a culture that treats engineering as craftsmanship, and learns from production.

That bar demands talent density: engineers who sharpened themselves in the world's best markets and chose to build here. Sophistication is the promise our products carry, so it has to be the standard our people set.

The Customer Interaction Company

Manifesto

One Platfo

Sawt connects a company's systems, runs its customer interactions across every channel, and continuously improves against business outcomes.

Every

Interactio

orm.

On one side of the platform are the systems where a business keeps its truth: the CRM, the ERP, the knowledge base, the customer record. On the other side are the places its customers talk: voice, WhatsApp, text, email. In between sits Sawt, the intelligent operating layer. It understands what the customer means, acts inside the company's systems, and learns from every conversation it carries.

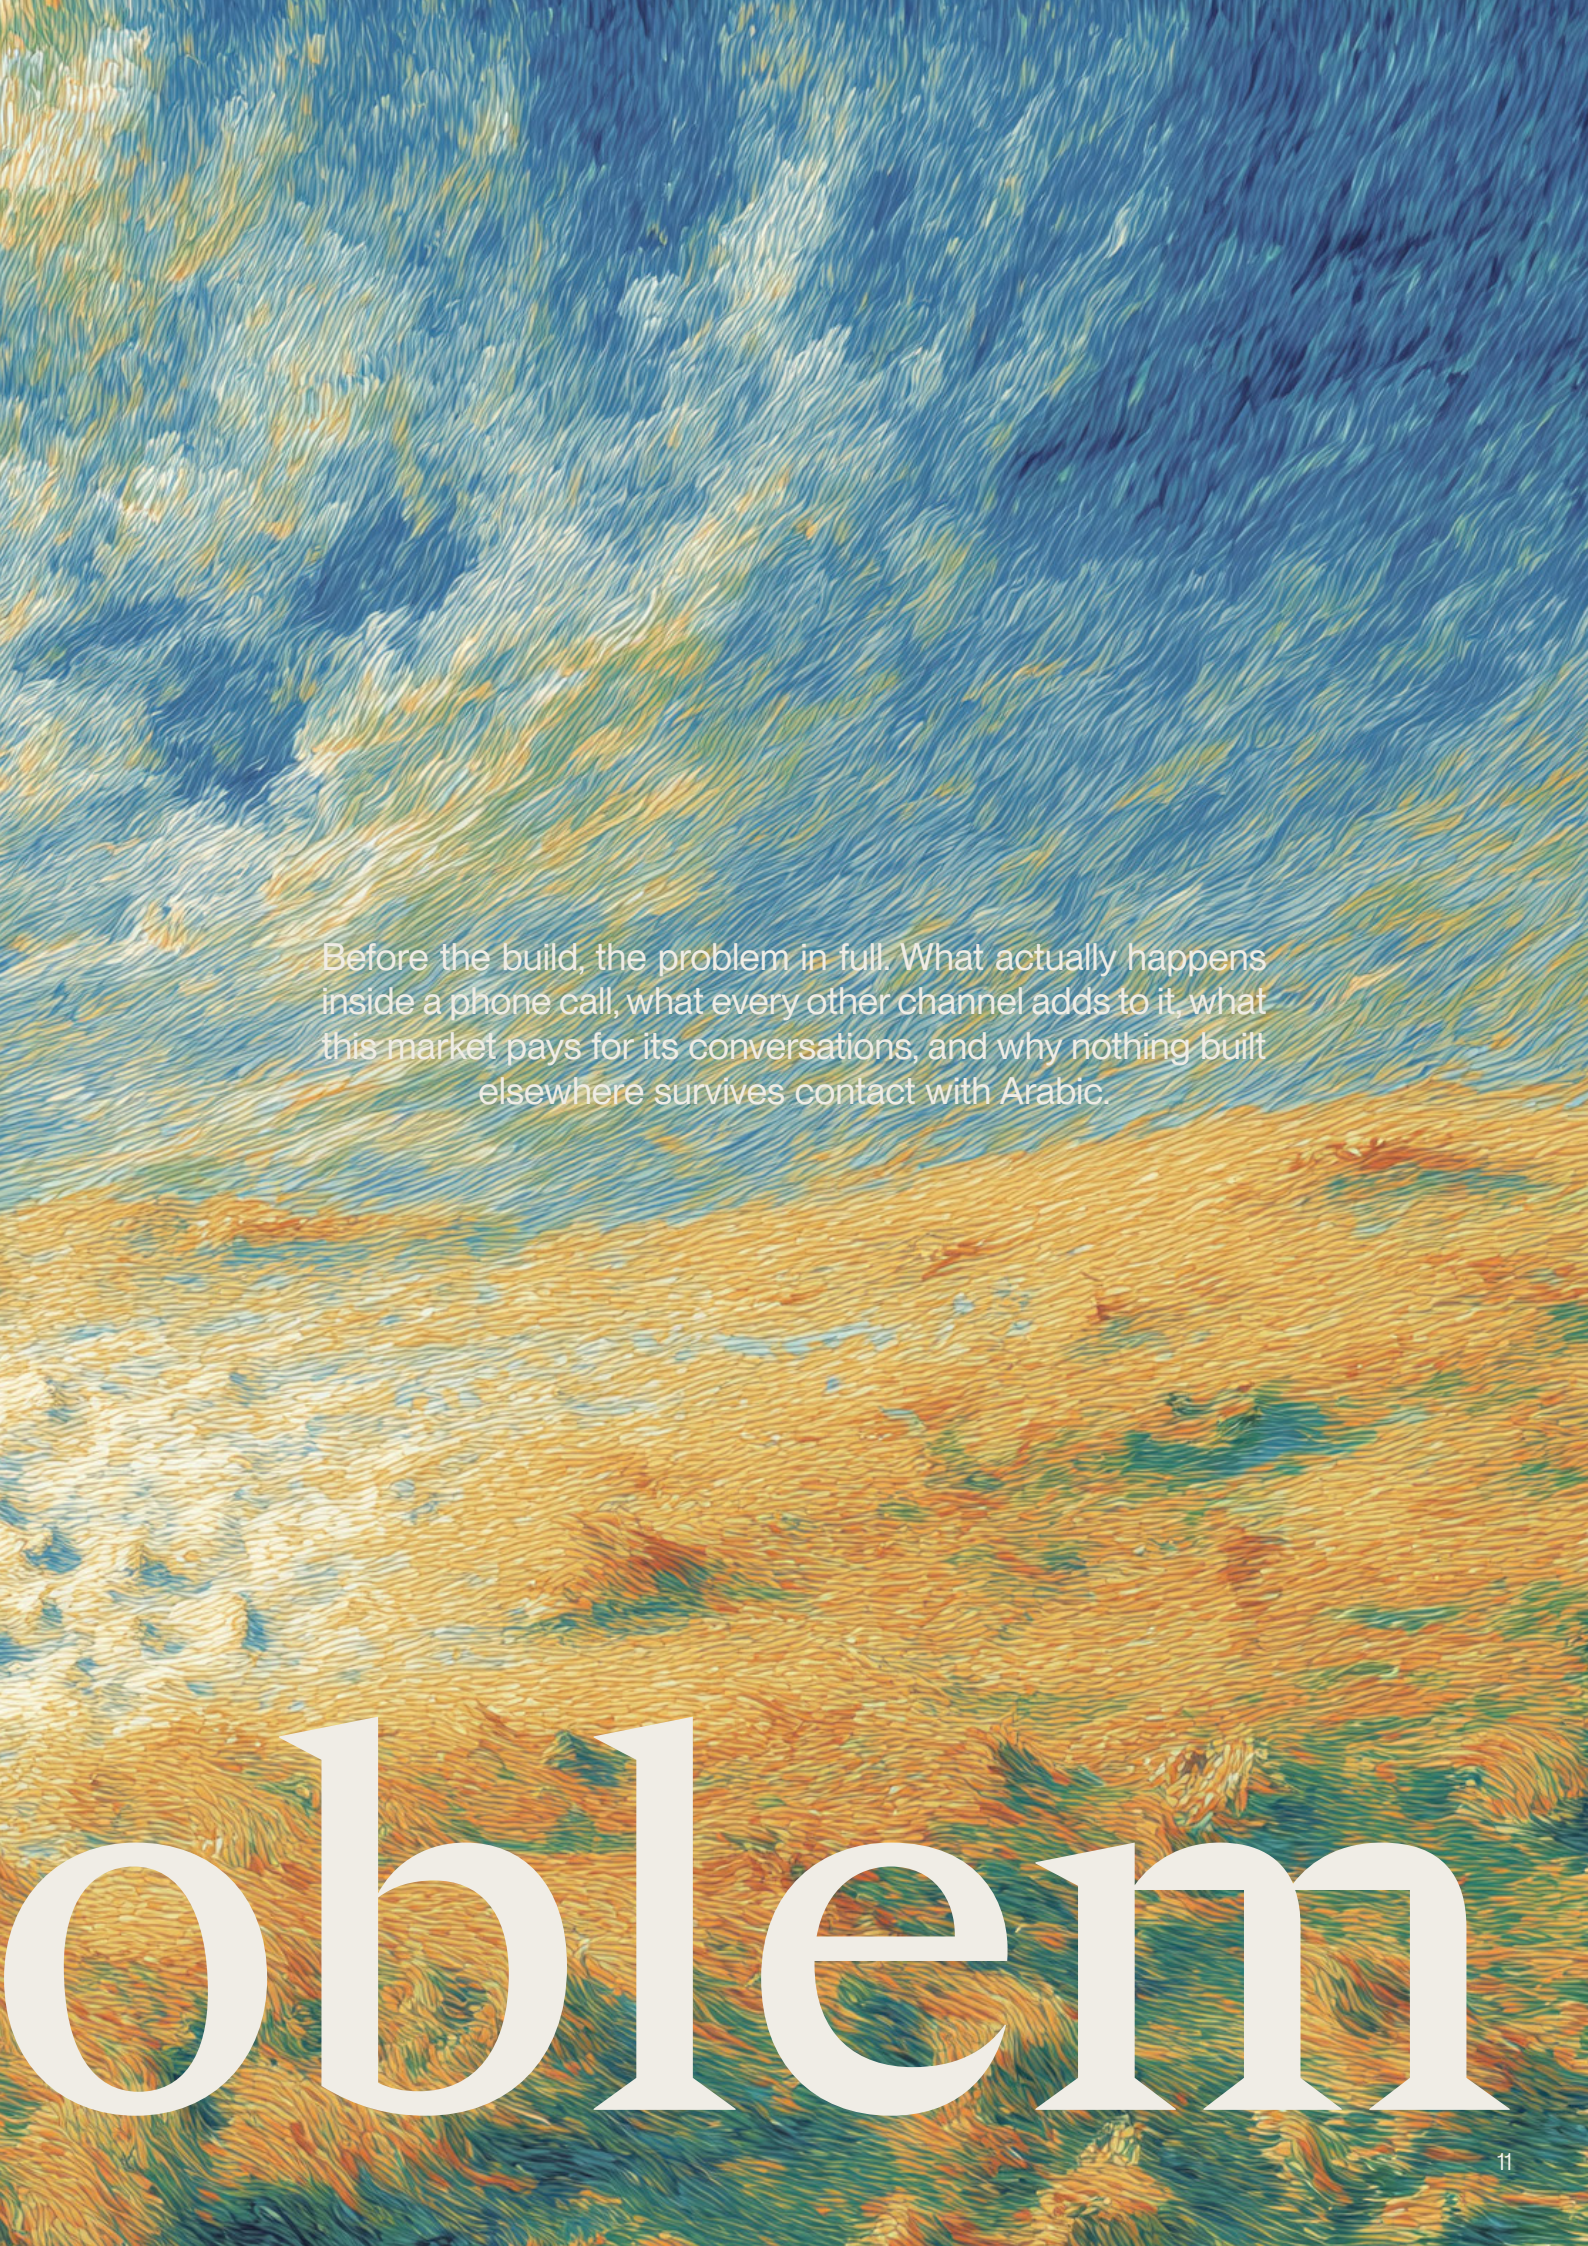
Customer

Sold as a platform. Delivered as a product. Measured in outcomes: conversations resolved, leads qualified, payments collected, customers converted.

n.

ACT I

The Pr



Before the build, the problem in full. What actually happens inside a phone call, what every other channel adds to it, what this market pays for its conversations, and why nothing built elsewhere survives contact with Arabic.

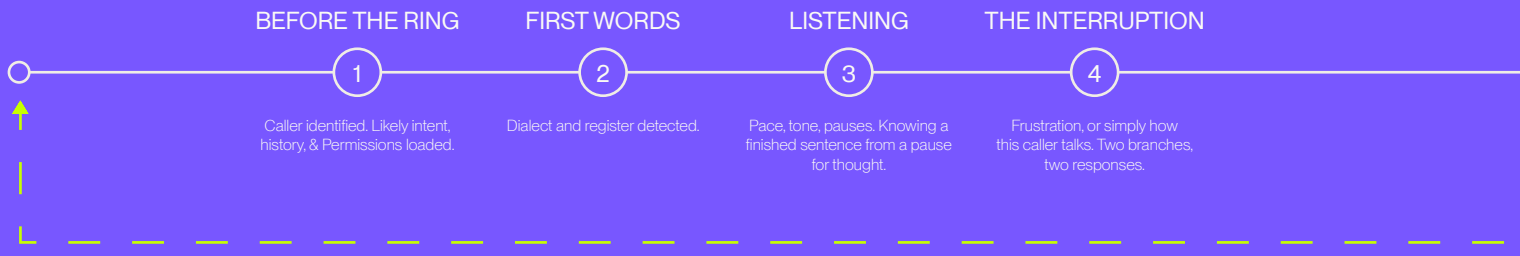
oblem

How Complex

A phone call looks simple from the outside. Someone dials. Someone answers. A conversation happens. Inside those seconds, a great deal is going on.

Before the call connects, the agent already needs to know who is calling, why they are probably calling, what their history looks like, and what it is allowed to say. The moment the line opens, the agent is listening. It listens to the words, and to everything the words arrive with: pace, tone, pauses, intent. It has to know when the caller has finished speaking and when they have simply stopped to think. It has to read an interruption correctly, because in one voice an interruption means frustration, and in another it is just how that person talks.

As the conversation moves, the agent makes decisions in milliseconds. Can it answer this on its own? Should it look something up? Send a link? Bring in a human? Each response has to land in the right dialect, at the right pace, with the right degree of formality. A mismatch in any one of them breaks the trust instantly, and trust is the whole conversation.



A Phone Call Is

Underneath all of this runs machinery that has to finish a full day's work inside a pause. The caller's words arrive as sound and have to become text. The text has to become meaning. The meaning has to be checked against policy, history, and the systems where the real answer lives. A response has to be composed, and the response has to become a voice, and every one of those steps has to complete in the space a person waits before wondering whether anyone is there. Fast and wrong breaks the trust. Right but slow breaks it too. And the caller can interrupt at any moment, and when they do, everything in flight is abandoned and the whole chain starts again, mid-sentence, without missing a beat. The conversation only feels natural while the machinery stays invisible, and the machinery only stays invisible while every stage runs close to instantly.

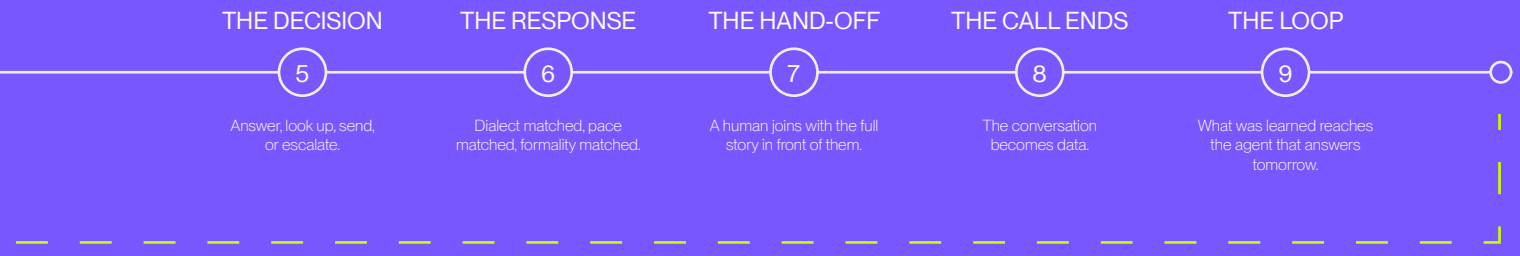
INSIDE THE PAUSE

A cutaway of the machinery between the caller finishing and the agent speaking.



Everything in flight is abandoned and begun again, mid-sentence, without missing a beat.

ALL OF IT , INSIDE A PAUSE THE CALLER NEVER NOTICES.



When the caller needs a person, the hand-off has to carry everything.
The human joins with the full story already in front of them, and the caller keeps going from exactly where they were.
Then the call ends, and the work carries on. The conversation becomes data. What went well, what went badly, where the agent hesitated, all of it feeds the next version of the agent that picks up tomorrow.
Most people never think about any of this. They just want the call to go well.

We think about nothing else.

The Whole Surface

A phone call is the hardest conversation software can have, and it is still only one of them.

Customers now reach a business every way people reach each other. They call. They type into a chat at midnight. They send a voice note from the car, a paragraph of frustration into an email, a single word into WhatsApp. Each channel has its own physics. A call happens in real time and punishes every hesitation. A chat runs on half attention, dropped and resumed over an afternoon. A voice note is a monologue that expects a listener. An email is formal, patient, and easy to ignore. Software that is fluent in one of them can be hopeless in another.

And customers refuse to stay in one. The same person starts in chat, gives up, and calls. She answers a survey by voice and follows up in writing. He confirms in WhatsApp what he asked on the phone. To the customer this is one conversation, moving through whatever door is nearest, and the business is expected to follow it without asking anything twice.

On the company's side, every one of those doors opens onto the same rooms: the same order, the same account, the same policies, the same systems where the truth lives. The channel is surface. The interaction is the thing itself.

So that is the unit we build around. Understand what the customer means, act inside the company's systems, learn from what happened, and hold the thread wherever the conversation goes next. One conversation, every channel, judged by whether it ended in an outcome.


The customer does not think in channels.
Neither should the software.

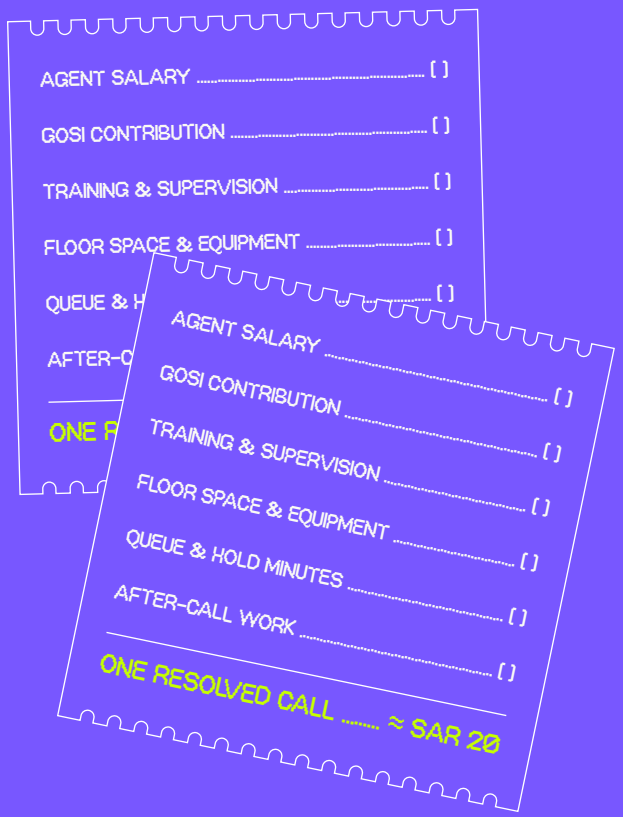
The Real Cost of a Call

Somewhere in Riyadh, a phone rings. A customer wants to move an appointment. The call takes four minutes, five with the paperwork, and by the time the file closes it has cost the business roughly twenty riyals. Multiply that by a million calls a year and the phone becomes one of the largest line items in the operation.

The arithmetic is harsher here than almost anywhere. Regulation keeps the work inside the Kingdom. Wages have climbed around 45 percent in five years. Staffing Arabic-capable teams around the clock is difficult, and with 97 percent of government services now digital, the volume keeps rising. A hundred riyals of customer service spend buys about three seats in the United States.
Here, the same hundred buys one.

Now ask what all that money buys the customer. Industry satisfaction sits near 78 percent. A single resolution takes six to eight minutes once hold time and wrap-up are counted. The story gets repeated to a second department, then a third. And still, 82 percent of Saudi consumers reach for the phone when the matter is sensitive, because when something truly matters, people want a voice. The most expensive channel is also the most trusted, and the most trusted is the one most likely to disappoint.

The people in the seats carry none of the blame. They care, and anyone who has watched them work knows it. They are running one of the hardest operations in business with tools built for other markets, in other languages, around other customers. The software barely understands the caller. The models behind it understand even less. The priciest service call in the world is answered through systems designed somewhere else, for somewhere else.



There is a quieter cost too. Talented people spend entire shifts on the same ten questions, asked a thousand different ways, while the work that actually needs a human waits. Software should carry that volume. Scale is what software does best, and when it carries the repetition, quality stops being rationed and people are freed for the conversations that deserve them. And the repetition arrives everywhere at once, on the phone, in the chat, in the voice note, in whatever channel the customer picked, which is why the answer has to live across all of them and care about the outcome of each conversation rather than the channel it came through.

Expensive, underserved, and trusted with what matters most.
No market in the world has more to gain from getting this right.

WHAT SAR 100 OF CUSTOMER SERVICE BUYS



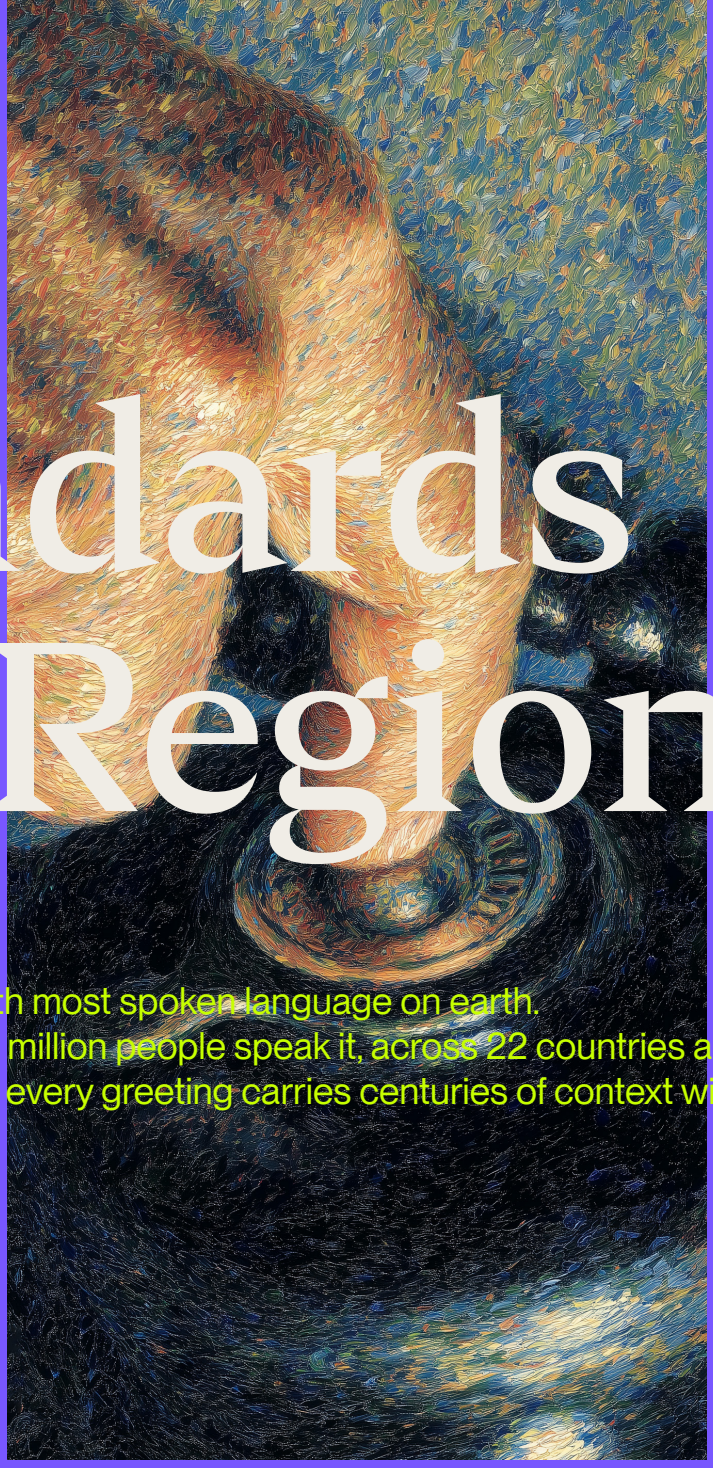
UNITED STATES ~3 SEATS (OFFSHORE-SERVED)



SAUDI ARABIA 1 SEAT (ONSHORE BY REGULATION)

The most expensive channel is also the most trusted, and the most trusted is the one most likely to disappoint.

CX Standards In The Region



Arabic is the fifth most spoken language on earth.
More than 400 million people speak it, across 22 countries and dozens of dialects, and every greeting carries centuries of context with it

Yet when an Arabic speaker calls a company, the system on the other end was almost always built somewhere else and trained on someone else's language. It cannot tell a Najdi accent from a Levantine one. It hears the words and misses the meaning.

The distance runs deeper than vocabulary. Arabic conversation shifts with tone, dialect, and expectation. A caller from Riyadh and a caller from Cairo can use the same word for opposite things. In one dialect, an interruption reads as aggression. In another, it reads as interest. A machine that serves this region has to carry all of that, fluently, at the speed of a live call.

Customers here hold the standard every customer holds. They expect speed, consistency, and an interaction that feels like it was made for them, in their language, on their terms. Being heard is the easy part. Being understood is the standard.

No global platform was ever going to close this distance.
It had to be built from the ground up. In Arabic, for Arabic, with Arabic speakers at the center.

400M+
Speakers

22
Countries

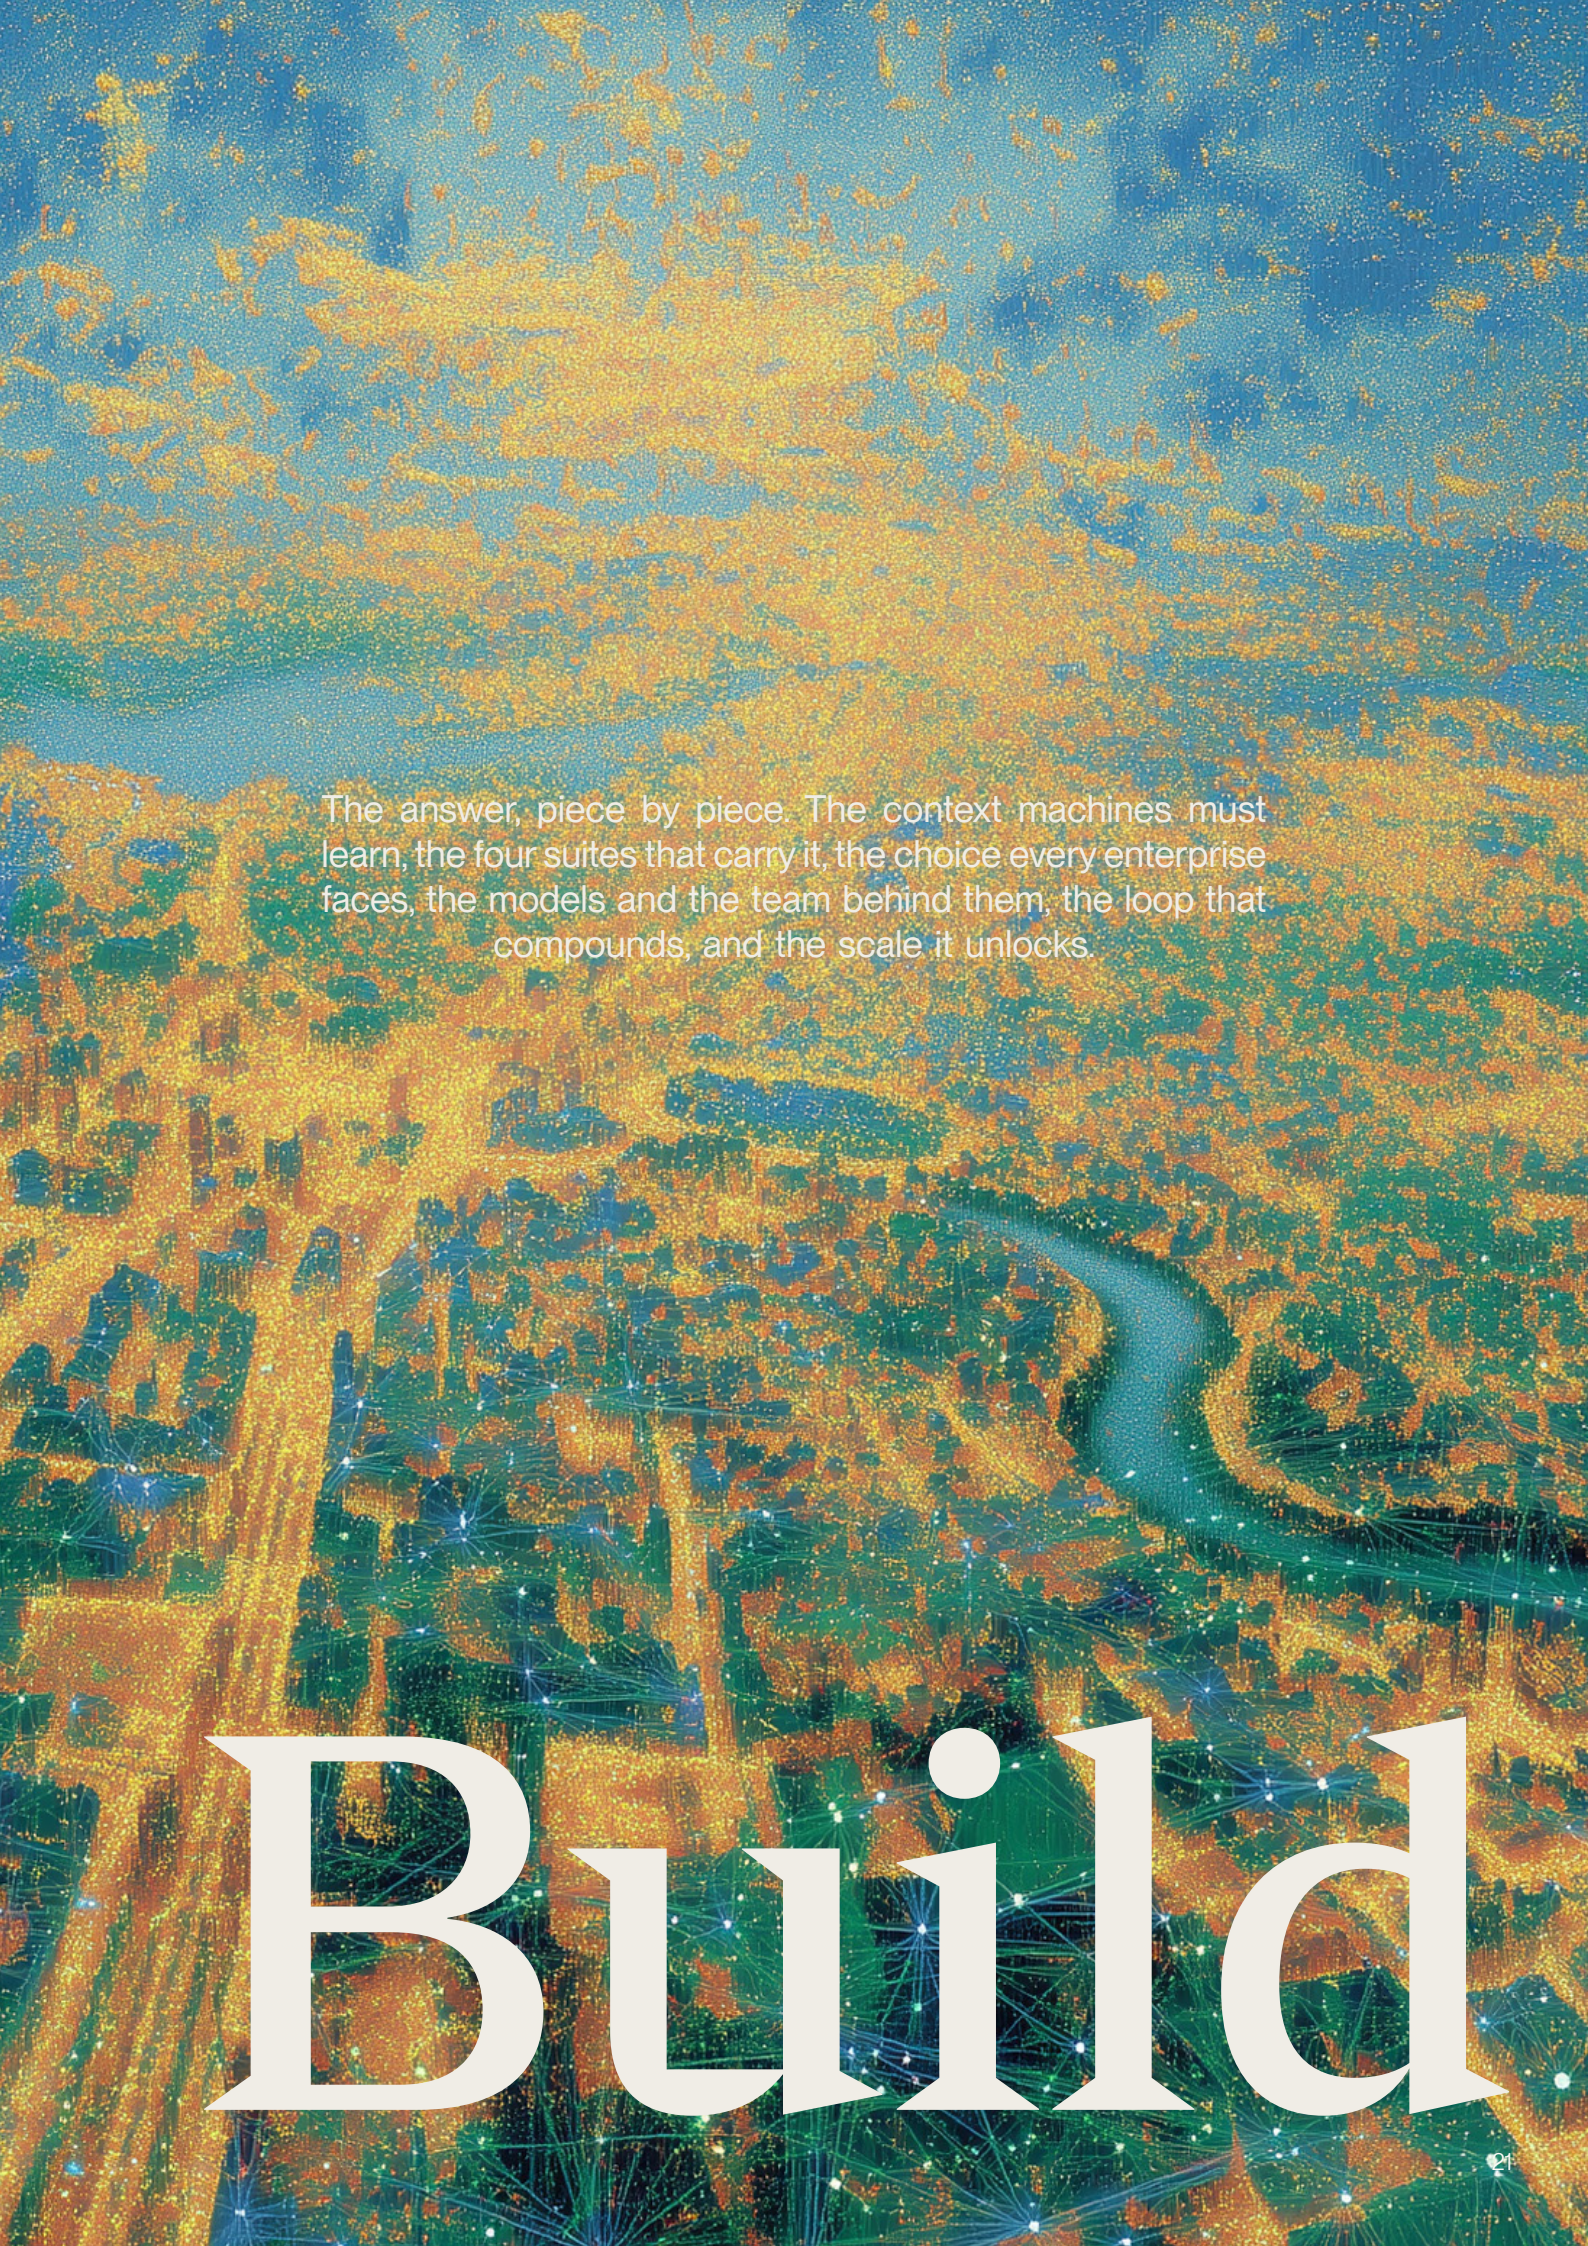
+30
Spoken Dialects

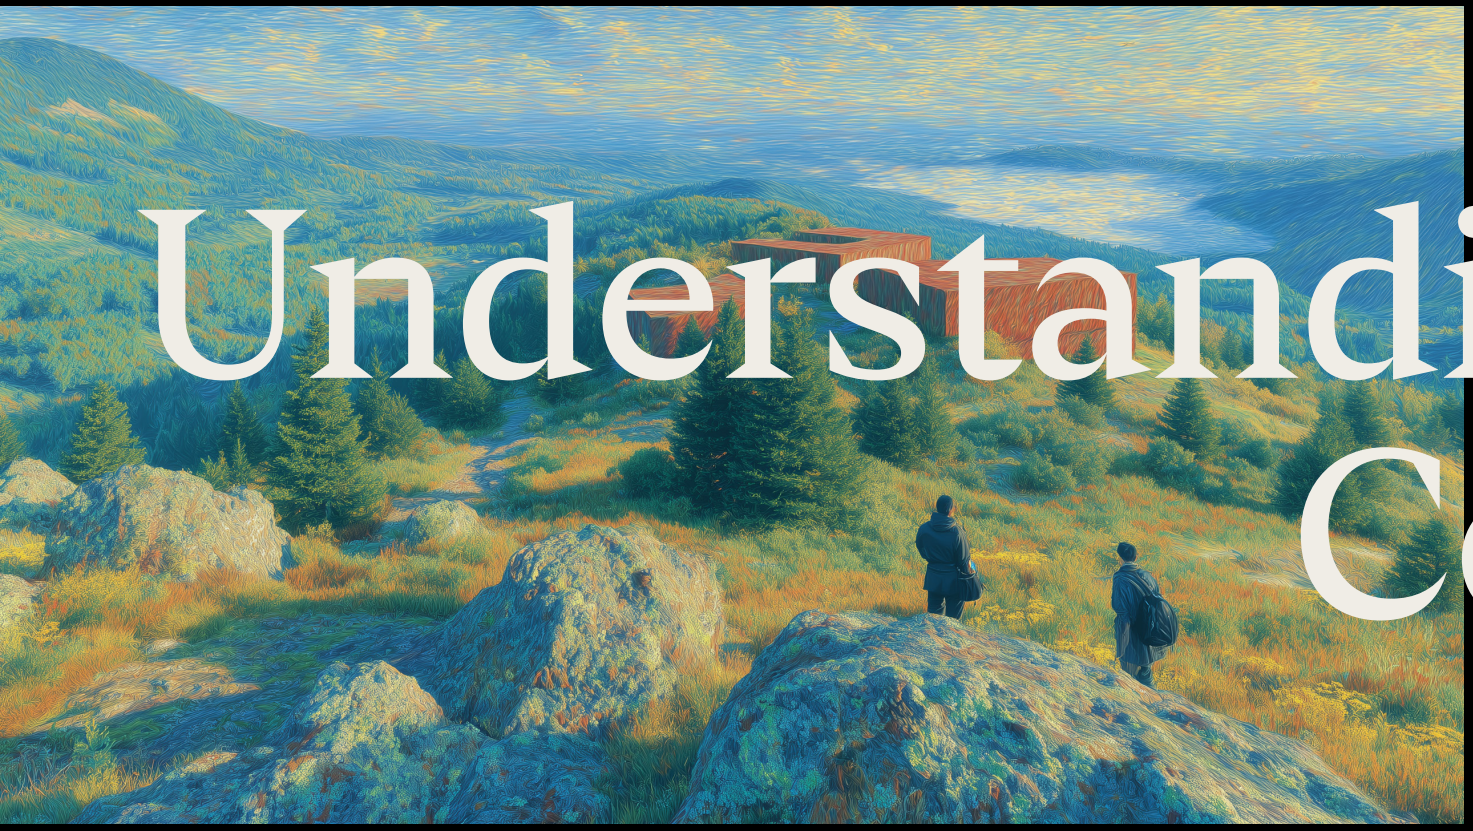
That is what Sawt was built to do.

ACT II

The




The answer, piece by piece. The context machines must learn, the four suites that carry it, the choice every enterprise faces, the models and the team behind them, the loop that compounds, and the scale it unlocks.
Build



Understanding Context

The best service anyone receives is still given by people who know them. The pharmacist greets you by name and asks about your father. The coffee shop starts your order when your car pulls in. The tailor remembers your measurements and your wedding date. None of this feels like service when it happens. It feels like being known.

What those people hold is context, and context runs far past language. It is knowing that when a caller asks about the papers, she means the contract from last week. It is knowing that a call after Isha is welcome in one household and an intrusion in another. It is hearing the difference between a customer who wants the short answer and a customer who called partly to talk. It is Ramadan hours, wedding-season rushes, the courtesies that open a conversation in Qassim and the different ones that open it in Jeddah.

A machine built elsewhere arrives with none of this. It can be taught the words and still miss every one of these signals, and the customer feels the miss instantly, the way anyone feels a stranger pretending familiarity.

Context has to be learned where it lives. It comes from the market, from real customers, from thousands of conversations in which the caller wanted things a manual never listed. Every one of those conversations taught our system something a textbook could never hold, and the lessons stay. The agent that answers in Riyadh today carries the accumulated context of every caller before.

We build for this region because we are of this region. The expectations of its customers are our specifications.

Language tells you what was said. Context tells you what was meant.

None of this feels like service when it happens. It feels like being known.

The 4 Suites

No business can predict every conversation its customers will have with it. You can write the policies, map the journeys, anticipate the questions and test the edge cases. Then a customer calls and does something nobody planned for.

That unpredictability has always made customer operations difficult to improve. A problem happens in one conversation, gets noticed in another, gets reported weeks later, and eventually reaches the people who can fix it. By then, thousands more customers have run into the same thing.
We think the distance between an interaction and what a business learns from it should be much shorter.

That belief shaped Sawt into four connected suites: Agent Creation, QA, Agent Operations and Insights. Four rooms of one house, arranged in a circle.

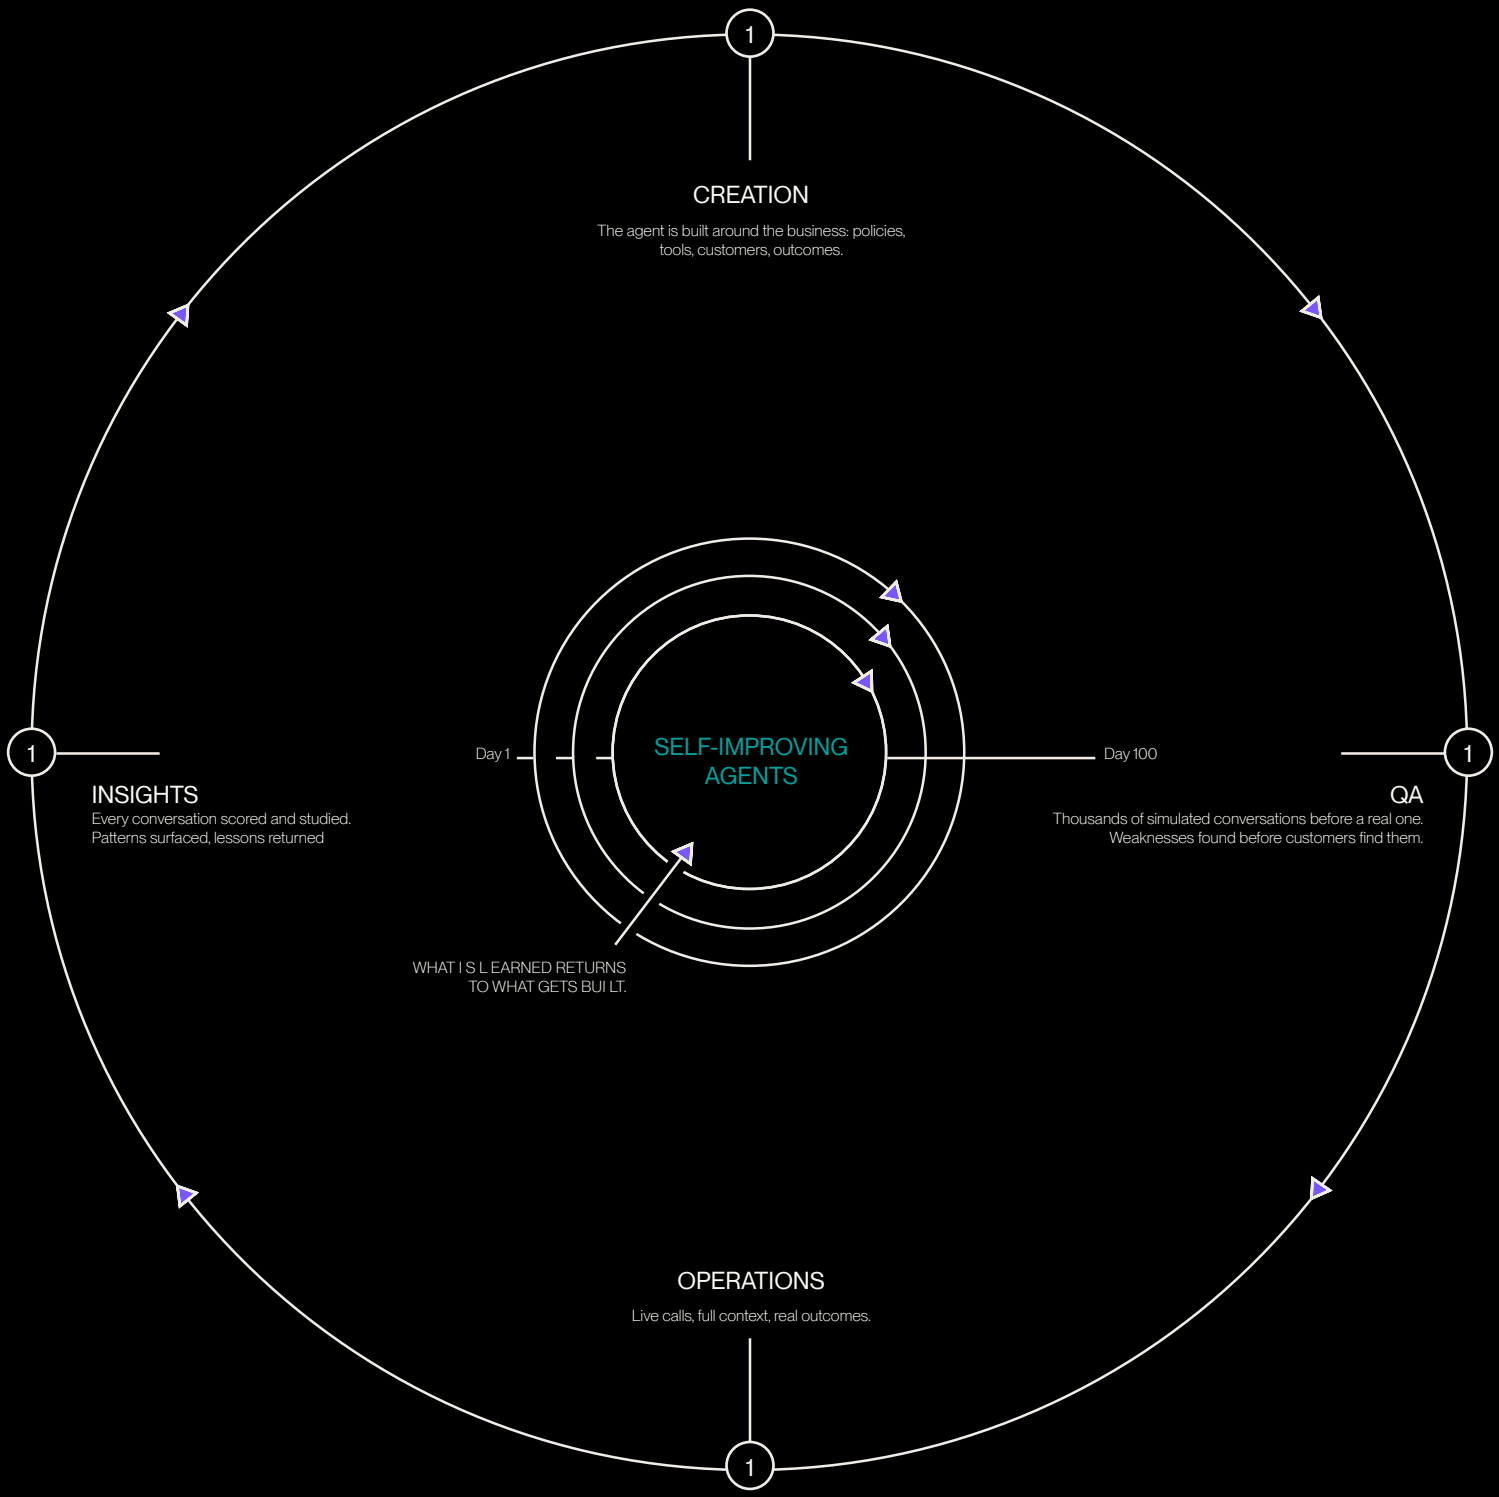
An agent is created around the business it represents: its policies, its procedures, its tools, its customers and its definition of a successful outcome. QA then challenges those decisions before production, simulating thousands of ways a conversation might unfold and finding the weaknesses before a customer does. Agent Operations takes what survives into the real world, live, at volume, on every channel a customer chooses: a call, a chat, a voice note. The suites carry any conversation a business runs, whatever the use case, because the loop cares about outcomes, and every use case has one.

Insights turns what happened across all those conversations into something the business can learn from. One customer struggling with a phrase is an exception. Thousands struggling with it is a pattern. A new objection, an unexpected behavior, a faster path to resolution: production reveals what planning misses, and what is learned returns to what gets built.

Creation gives the agent a starting point. QA challenges it. Operations gives it experience. Insights gives that experience meaning. Then the circle closes.

The first version of an agent is built from what the business already knows. Every version after carries what the business has learned from its customers.

The agent on day one has instructions. The agent on day one hundred has experience.



The agent on day one has instructions. The agent on day one hundred has experience.

Buy or in-house?

Every enterprise that decides to put AI into its customer conversations reaches the same fork. Build in-house on top of a raw model, or build on a platform made for this.

The strongest AI companies in the world have written honestly about what that choice involves, and their account matches ours. The demo is cheap and the production system is expensive. A capable model gets a prototype talking in a week. Then come the parts the prototype never showed. Testing against thousands of ways a conversation can go wrong. Connecting the systems where the customer's answer actually lives. Guardrails around what the agent may promise, say, and do. Watching every conversation, catching every failure, and shipping the fix before the failure repeats. Rebuilding it all when the next model generation arrives, which it does every few months. The build is a permanent engineering commitment wearing the costume of a project.

A platform changes what each new instruction costs. When the system already

WHAT THE DEMO SHOWS

A talking agent

A good first call

knows how to verify a caller, consult a policy, reach into a backend, and hand off with context, the next command builds on all of it. That is how agents graduate from answering a question to carrying real work over a long horizon, tasks that span days, systems, and follow-ups. Long-horizon agents are a platform achievement. No single build gets there alone.
And the platform only matters once it is embedded. This is the most important thing. An agent bolted onto the surface of a company answers questions. An agent embedded inside it, wired into the channels customers actually use and the systems where the work actually happens, completes the work itself. Embedding is where the value lives, and embedding is operations: patient, unglamorous, and learned only by doing it inside real enterprises.

Something larger sits underneath. Software is changing shape. For twenty years it sold features, priced by the seat, and left the outcome as the customer's problem. It is turning into a service that sells the outcome itself: the resolved conversation, the confirmed address, the collected payment. Businesses will compare providers on outcomes delivered and choose accordingly, and feature lists will stop deciding anything. What remains is an applied machine learning and operations game. That is the game we chose, in the market where it matters most.

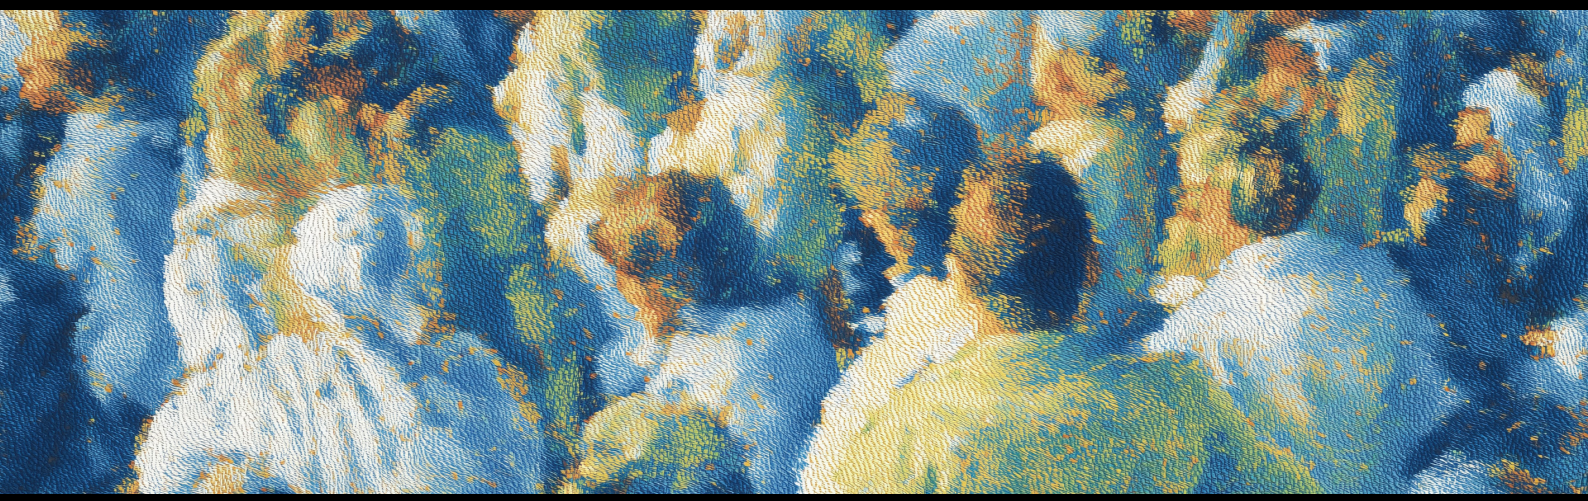
Build what makes you singular. For the rest, choose partners who build nothing else.

WHAT PRODUCTION DEMANDS

Testing at scale
Integrations
Monitoring
Dialect coverage
Hand-offs
Model migrations
Around the clock reliability

Long-horizon agents are a platform achievement.

The Best Arabic Models,



& The Team Behind Them

Building an AI agent that speaks Arabic is a different task from translating an English one. Arabic has more than thirty spoken dialects. Words change meaning between regions. Pronunciation shifts between cities. A model trained on Modern Standard Arabic will fail a caller from Jeddah the way a model trained on textbook English would fail a caller from Glasgow.

So we built our own models, from the ground up, trained on the way Arabic is actually spoken. Speech recognition that holds up on a real phone line, with its background noise, its overlapping voices, its half-finished sentences. Speech synthesis that sounds like a person from here, in the right dialect, at the right cadence. A language model that reads the meaning under the words, follows policy, and answers the way the situation calls for. All of it runs in-house, and all of it can live inside a client's own infrastructure.

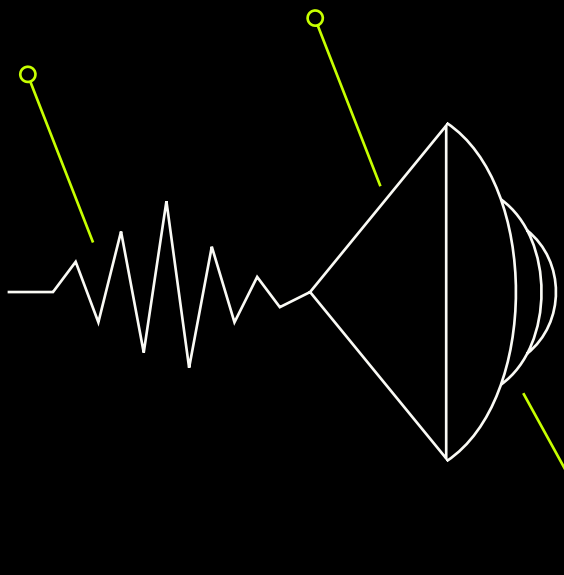
Behind the models is a team, and the team has a rule: nothing ships on a feeling. Every model, every feature, every change is measured against real Saudi conversations before it reaches a customer. The test sets are built from the hardest material production can offer, the noisy lines and the thick dialects and the calls where the first version stumbled. When a model clears the bar, it ships. When it misses, the miss becomes training material. The bar itself rises every quarter, because production keeps supplying harder questions.

That discipline produced a result the team is quietly proud of. On the conditions that matter here, real Saudi callers on real phone lines, our models outperform every system we have been able to test against them, including those built by the largest companies in the world. Scale reads the whole internet. It has never worked a phone line in Qassim. A model taught by this market, on this market, beats a giant that has never been here.

The engineers who built this sharpened themselves in the world's best labs and markets, and chose to come build here. Ask them why and the answers vary, but one thread repeats: nowhere else does the work matter this much to the place it serves.

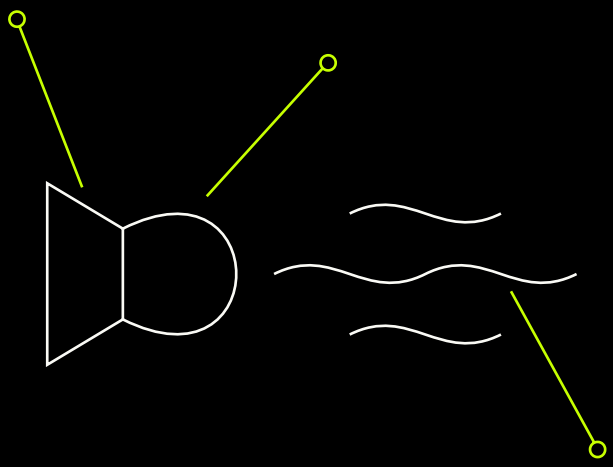
Our models speak Arabic the way the Kingdom speaks it.

SPEECH RECOGNITION (ASR)



hears dialect, noise, interruption.

SPEECH SYNTHESIS (TTS)



a voice from here, at conversational cadence

Scale reads the whole internet.
It has never worked a phone line in Qassim

EXHIBIT MEASURED IN THE BLIND

Nothing here ships on a feeling. Blind rankings, real Saudi phone lines, and production failures turned into test material. These are the standings.

SPEECH RECOGNITION (ASR)

Public Benchmarks — Lower is Better

SAWT Telephonic Benchmark

10.50

WER - CER 4.90

16.56

Best Competitor - WER

Clips Under 2 Seconds

11.75

WER

22.54

Best Competitor - WER

Sada

23.98

WER - CER 10.43

28.52

Best Competitor - WER

MGB-2

9.33

WER - CER 3.52

10.16

Best Competitor - WER

23.98 Entity Recall

16.8% Word Error Rate

8.3% Character Error Rate

SAWT Internal Telephonec Benchmark

Rescored 23 August 2026

Entity Recall on Real Phone Calls — Higher is Better



A transactional call fails on the entity: the name, the amount, the ID. This is the metric that decides whether a call completes.

SPEECH SYNTHESIS (TTS)

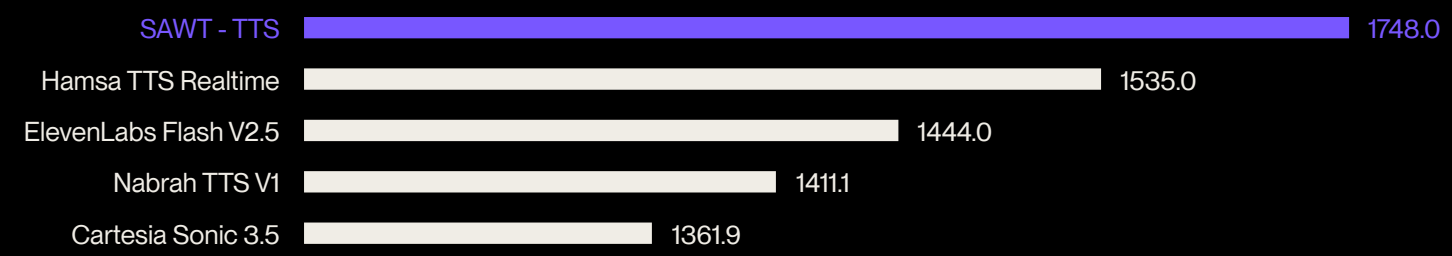
61.5%

Ranked First in Blind Evaluations

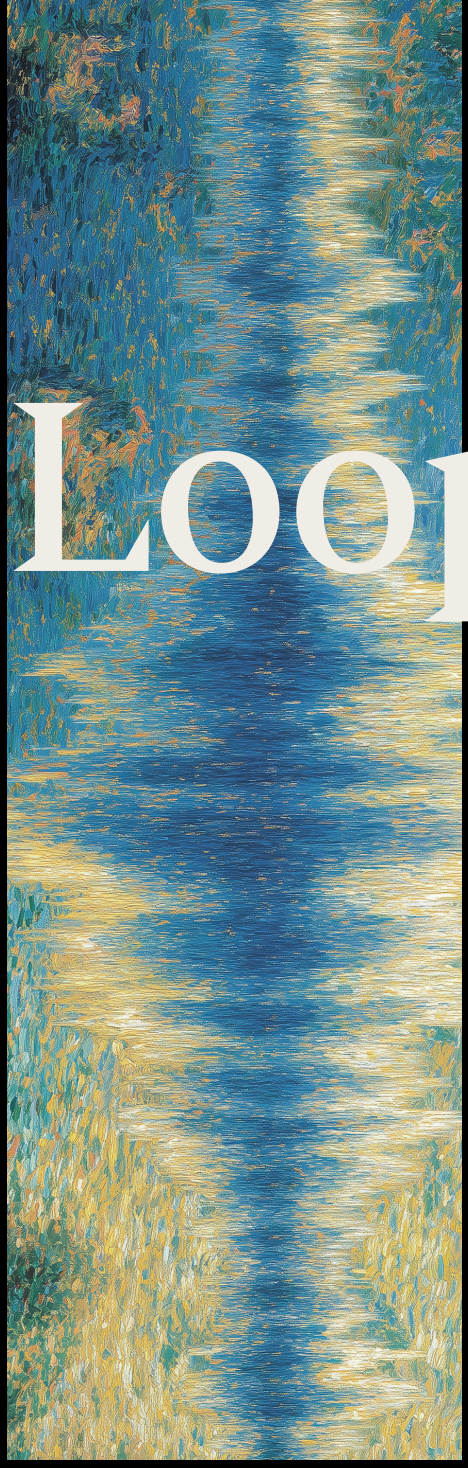
Public TTS Arena Leaderboard — August 2026

System	ELO	Average Rank (Lower is Better)
SAWT - TTS	1748.0	1.615
Hamsa TTS Realtime	1535.0	3.038
ElevenLabs Flash v2.5	1444.0	3.308
Nabrah TTS v1	1411.1	3.346
Cartesia Sonic 3.5	1361.9	3.692

Elo Ratings - The Same Five Systems



The Signal Loop



Writing software has never been easier. The tools are open, the models are public, and code that once took a team a year now takes a good engineer a month. Whatever can be written quickly can be copied quickly. So the question that decides who wins has moved. The question is what a company owns that cannot be written.

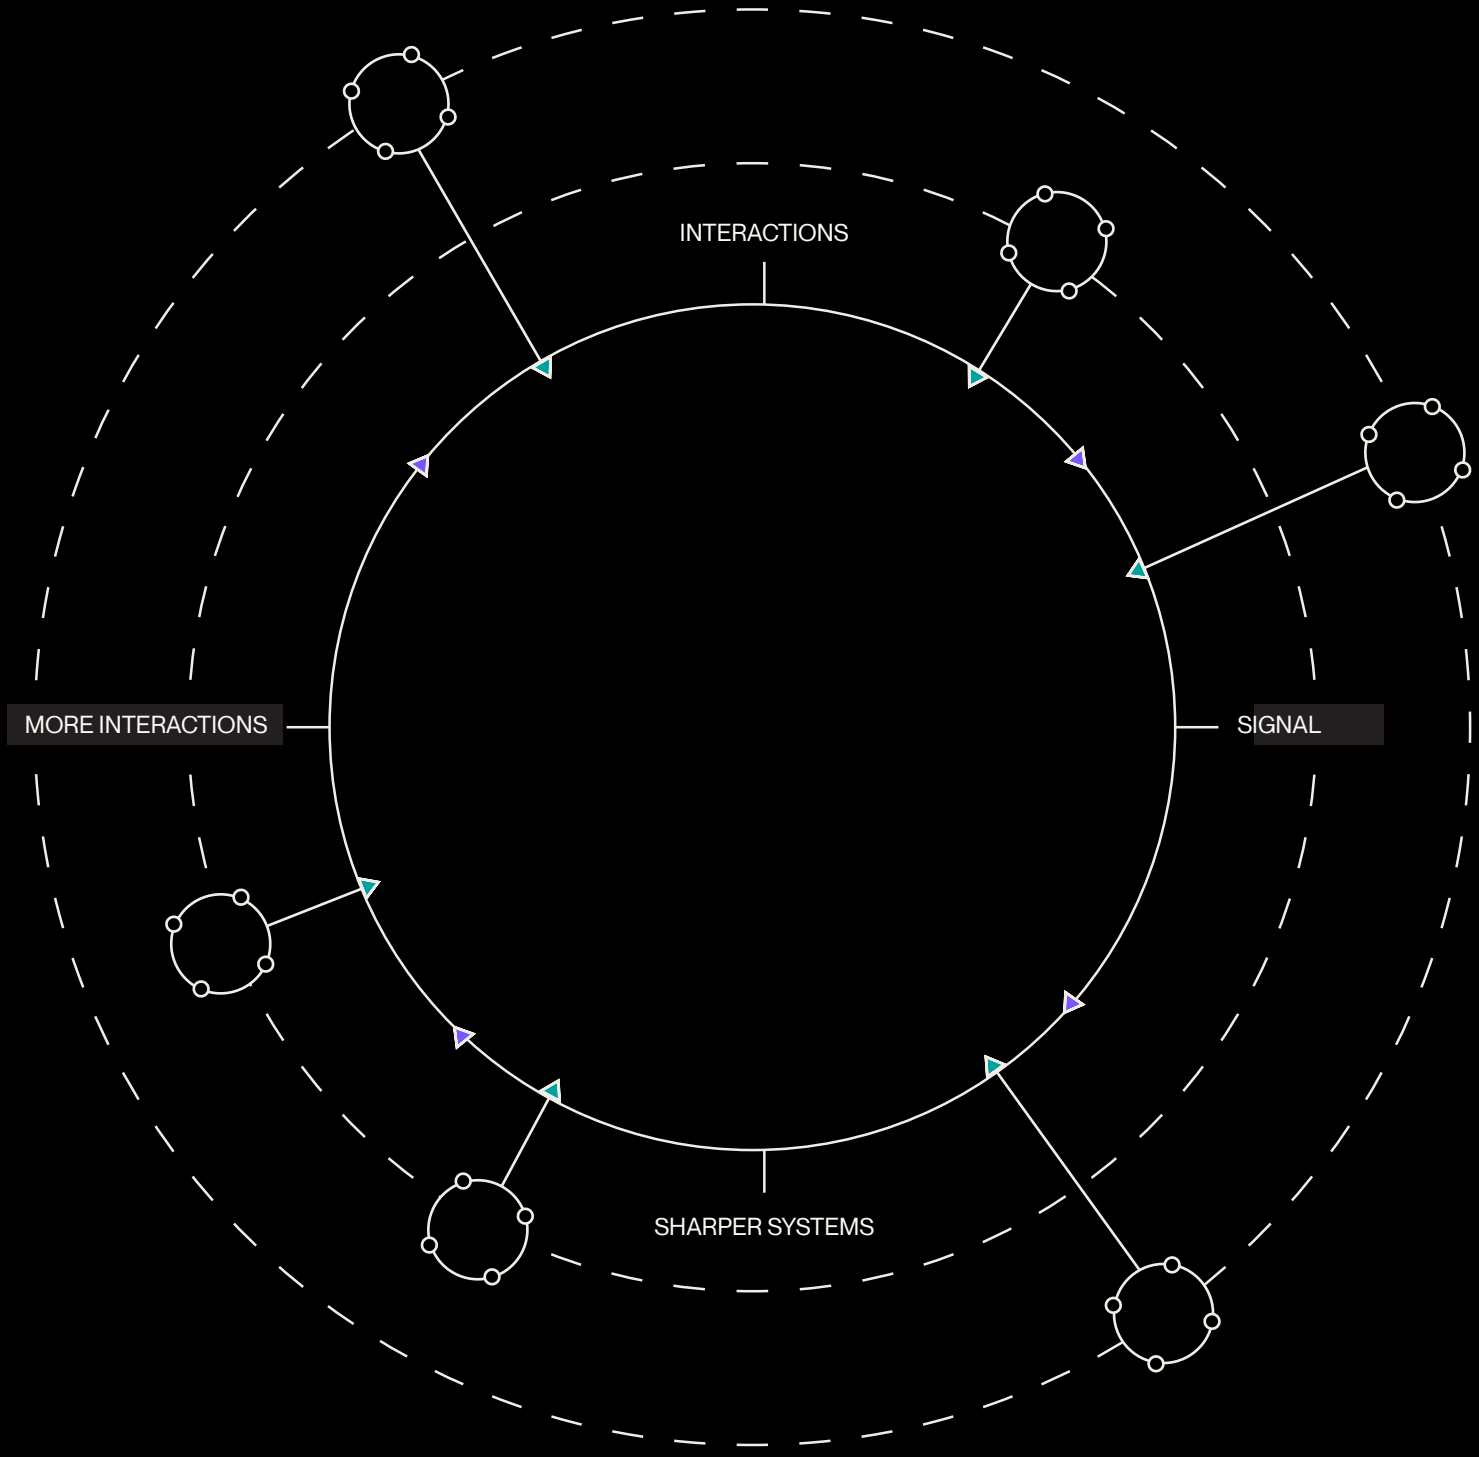
Ours is signal.

Every conversation our agents handle leaves something behind. The phrasing that confused a caller in Dammam. The pause that meant hesitation and the pause that meant the caller was driving. The objection nobody scripted. The shortcut to resolution that no one designed, discovered live, on a Tuesday, in dialect. Millions of these fragments, each small, each worthless alone, together form a record of how this market actually talks that exists nowhere else on earth.

The loop is what turns the record into an edge. Production surfaces the moments where the system struggled. Those moments become the exact material the next version trains against. The improved version returns to production sharper, meets harder cases, and sends those back too. Around and around, every day, across every deployment at once.

This is why time deepens the advantage. A competitor can match a feature in a quarter. To match the signal, they would need the years of live conversations that produced it, and those years are already ours. Each new deployment starts ahead of the last, carrying everything every earlier deployment learned.

The circle from the Four Suites is drawn small, around one agent. Draw it large, around the whole company, and it becomes the engine of the business: interactions produce signal, signal produces better systems, better systems win more interactions. Software is written. Judgment is earned.



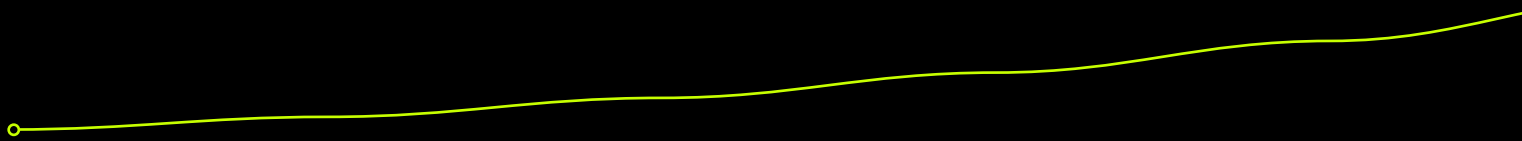
Software is written. Judgment is earned.

Scalability

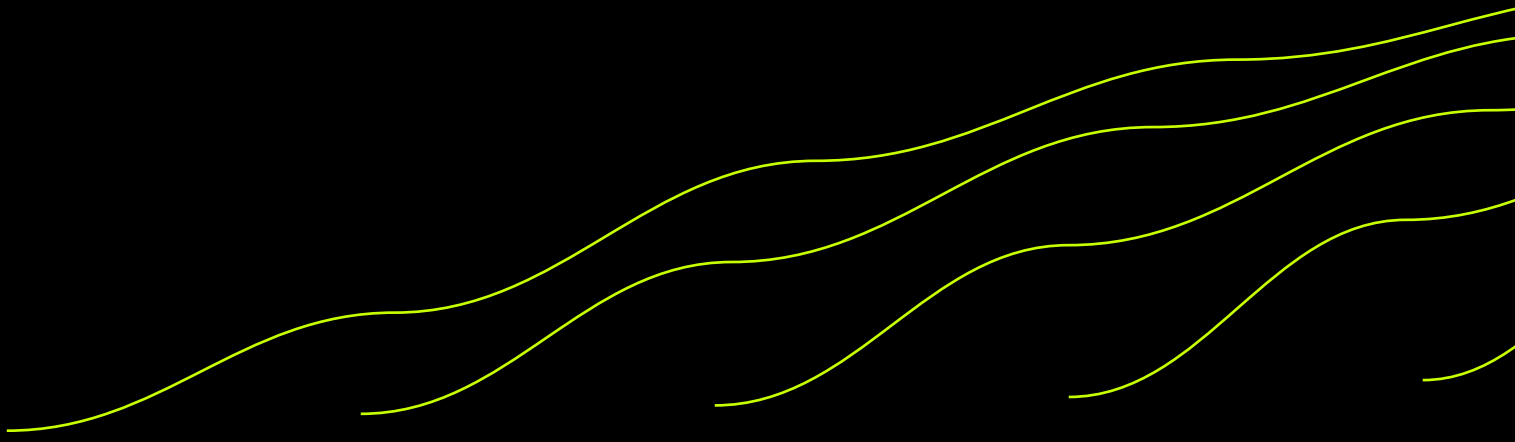
It is easy to make one interaction memorable. Listen closely, remember the context, understand what someone needs, and take care over the details. Small businesses do this naturally. The pharmacist knows your name. The restaurant remembers your order. The person answering the phone knows why you called yesterday. Scale makes that harder.

As thousands of customers become millions, the operation grows between the business and the person it serves. More queues, more systems, more people, more handovers. Demand arrives unevenly and sometimes all at once. What was once personal becomes a process, and somewhere along the way, quality begins to get rationed.

END TO END WITH INTERACTIONS - WITH TIME



USE CASE METRICS - WITH TIME



The customer never experiences that complexity. They experience one interaction.

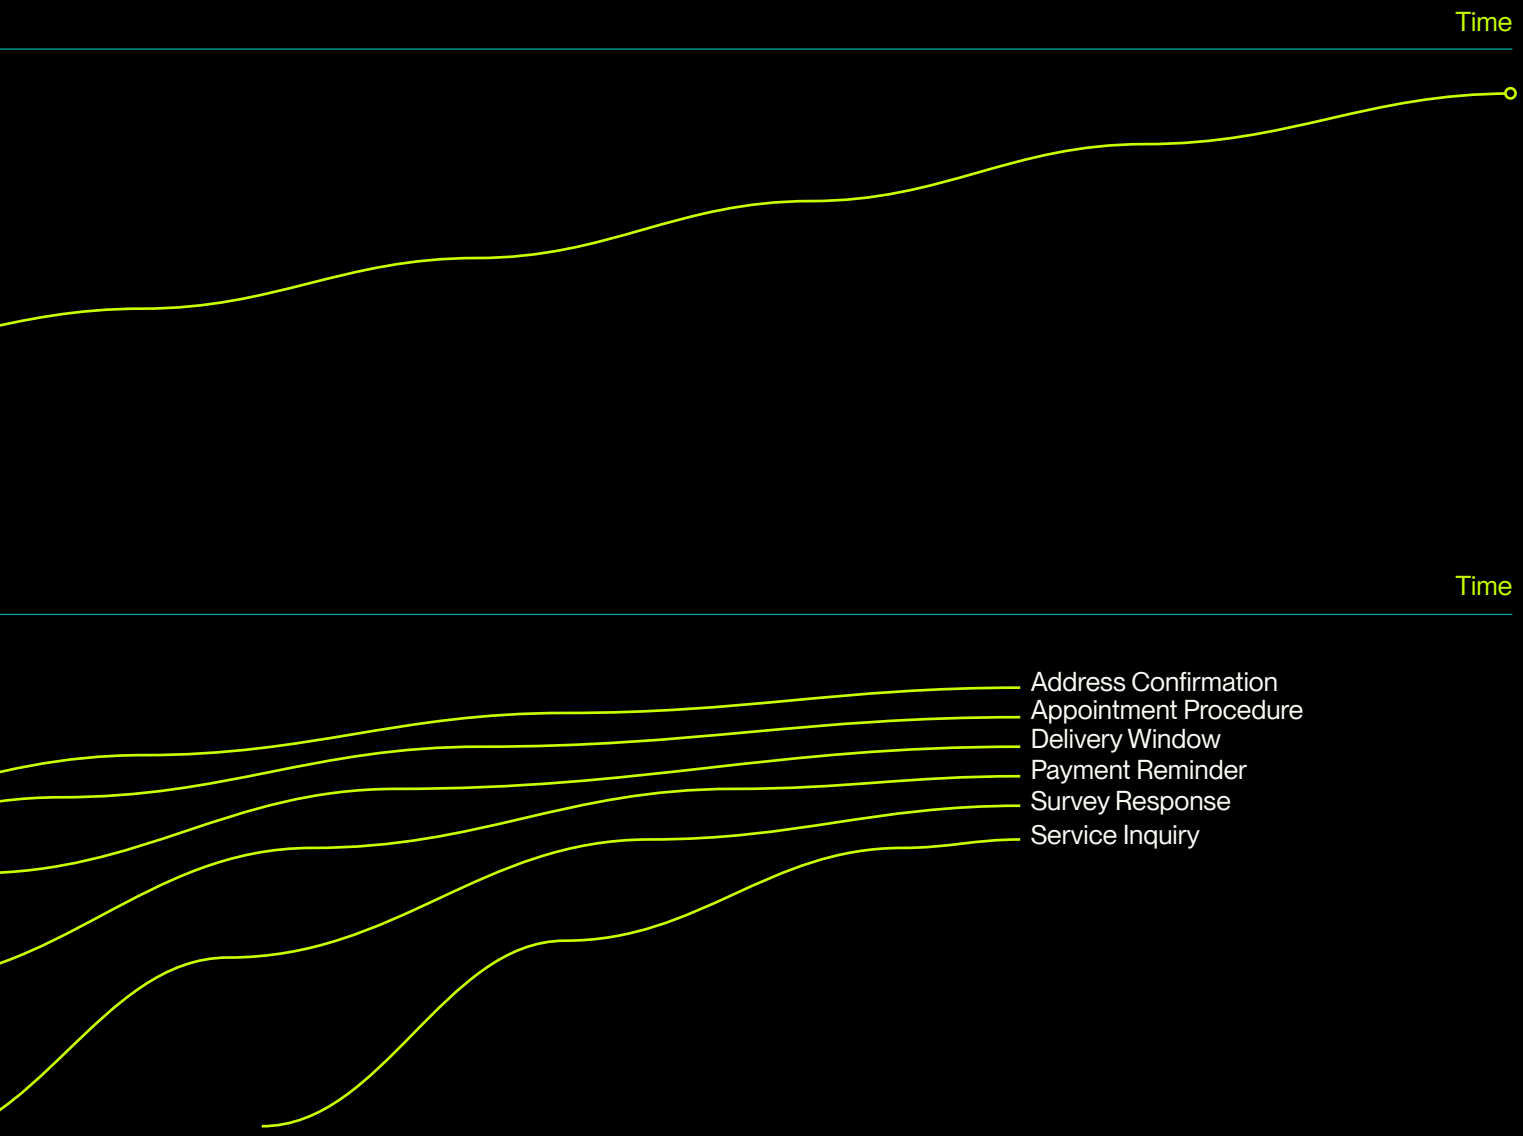
They still expect the business to know who they are, understand why they reached out, speak their language and resolve what they came for. Whether ten or ten thousand other customers are asking for the same thing should make no difference to them. That is the real challenge of scale: the numbers grow, and the standard has to hold.

Software changes what is possible here. Thousands of interactions can happen at once, each carrying a different history, dialect, intent and context, while every one of them stays personal.

Scale also creates something small businesses never had: the ability to learn from every interaction at once. One customer teaches you a preference. Millions reveal patterns. Where people hesitate. What they repeat. Which conversations resolve and which ones break. Each interaction leaves behind context that makes the system sharper for the next.

For most of business history, growth forced a choice between the intimacy of knowing your customers and the efficiency of serving millions of them. We think that choice is disappearing.

Making one customer feel understood is good service. Making millions feel understood is infrastructure.

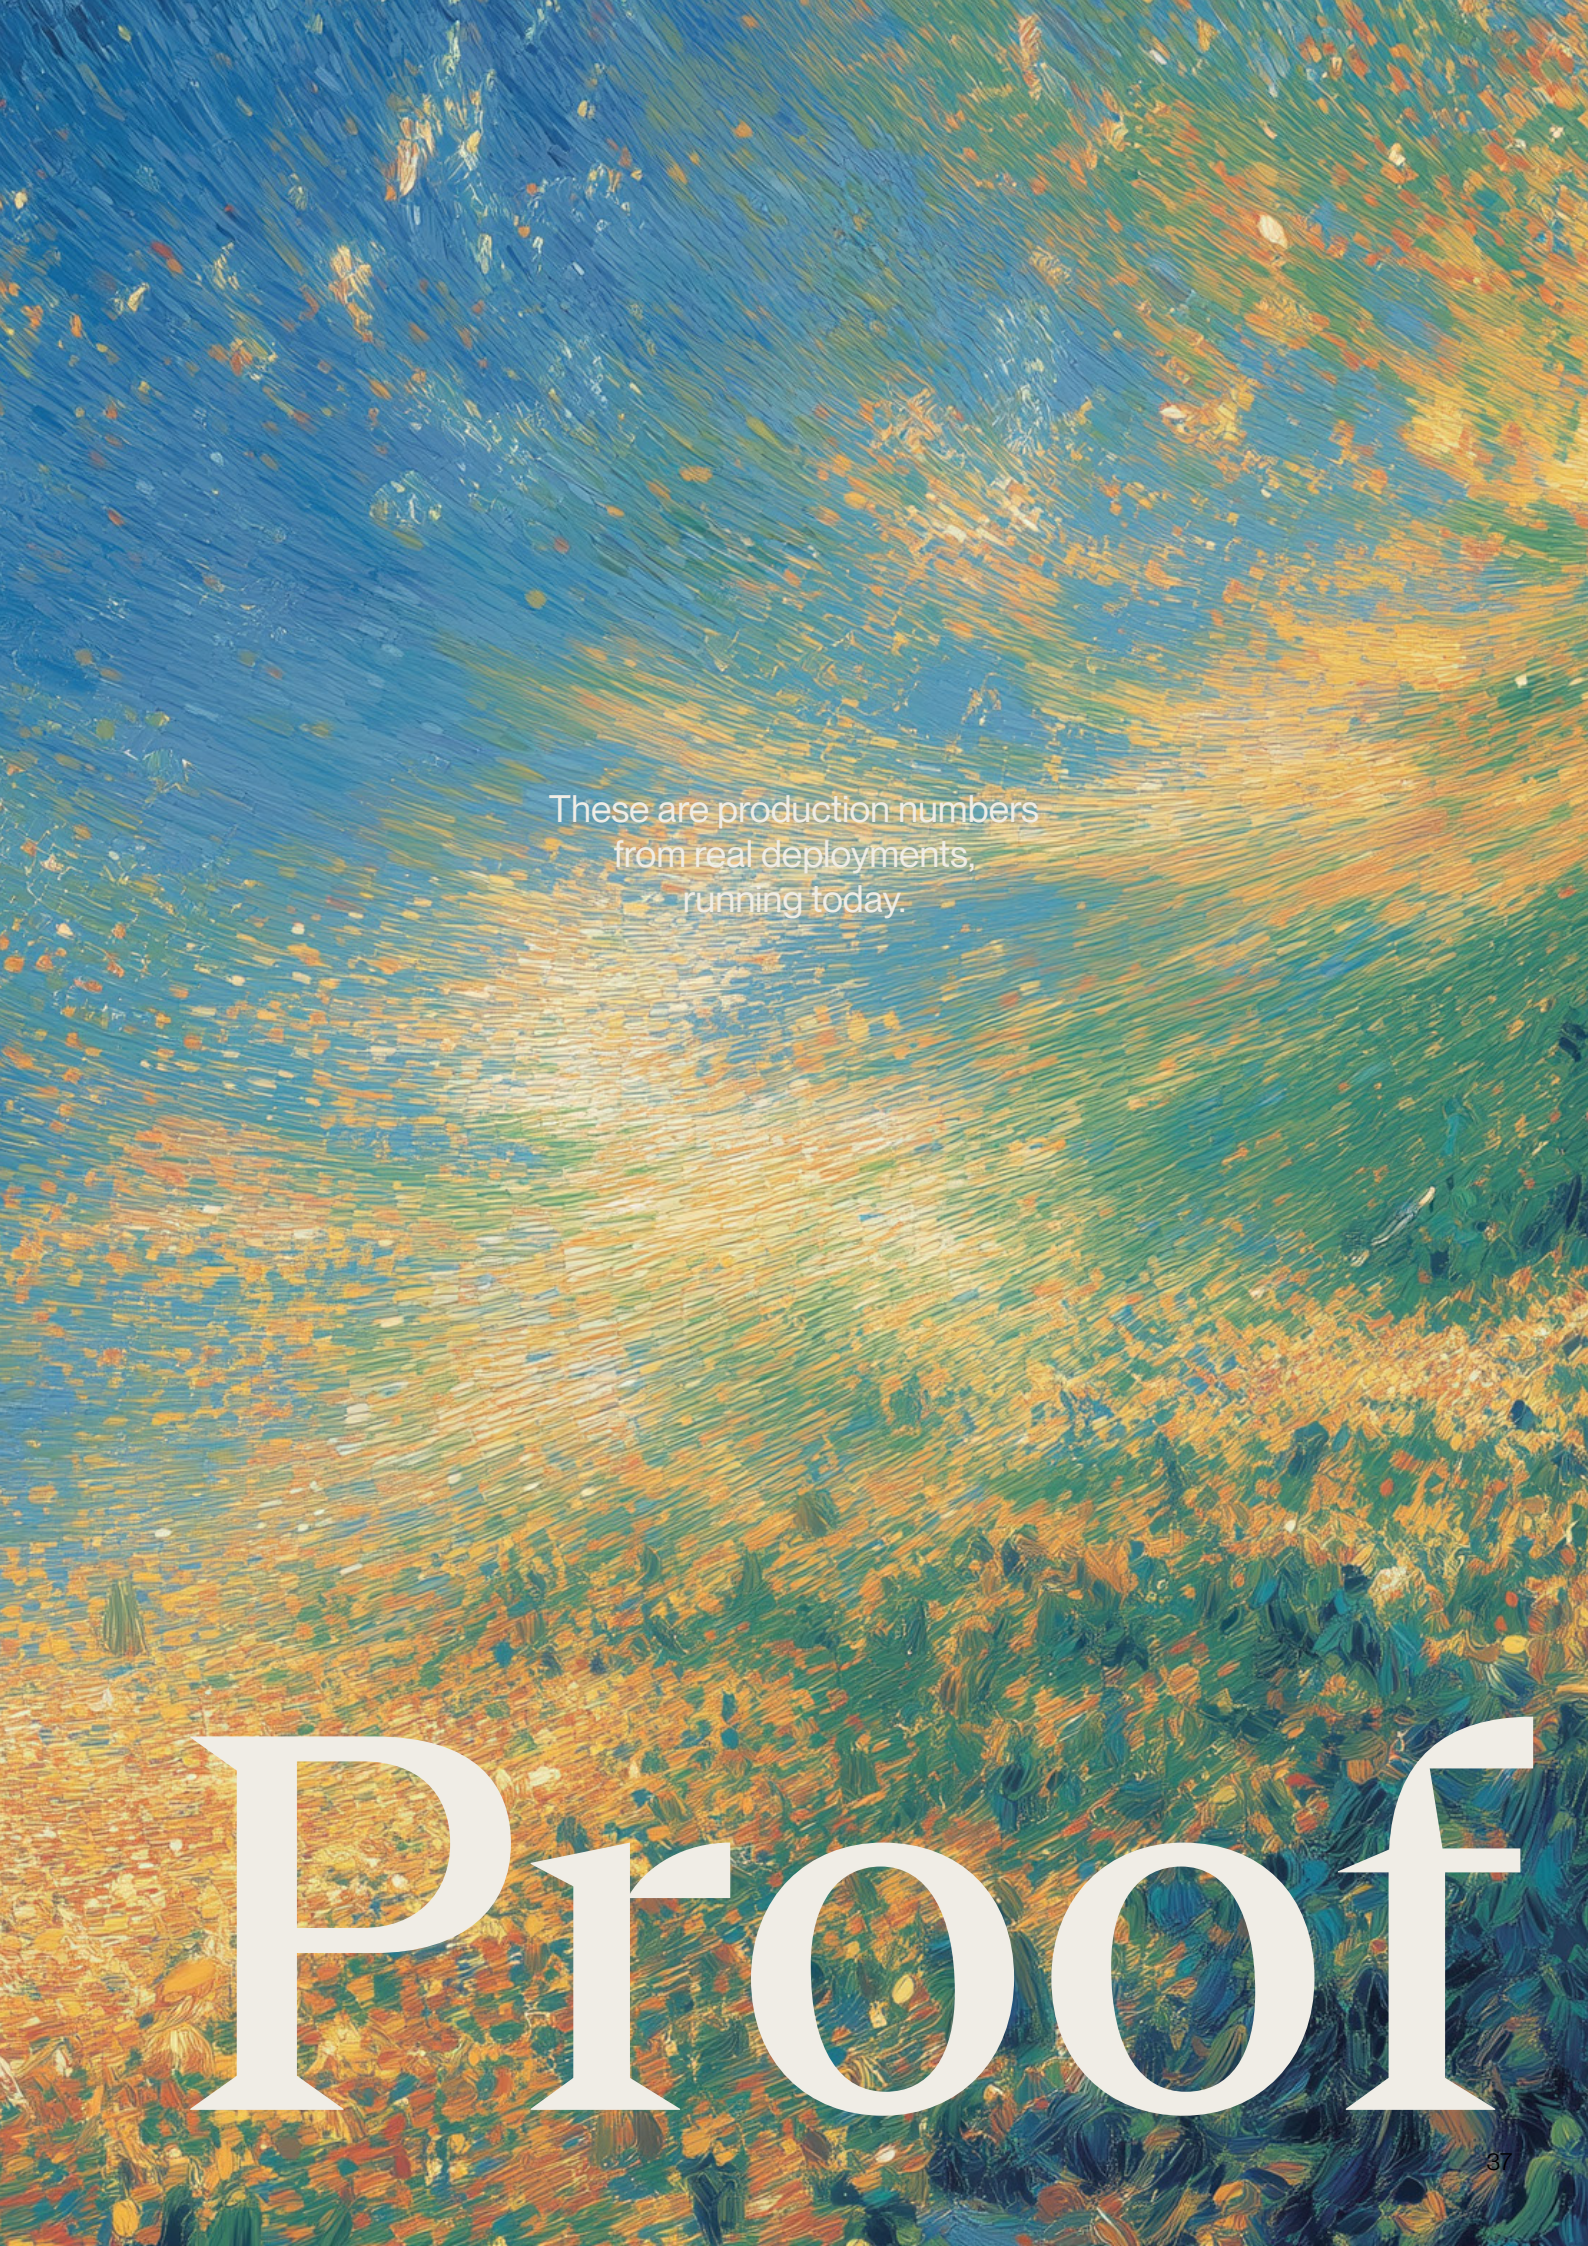


Making one customer feel understood is good service.
Making millions feel understood is infrastructure.

ACT III



The


These are production numbers
from real deployments,
running today.
Proof



COST REDUCTION - ONE CONVERSATION - AUTOMATED END TO END

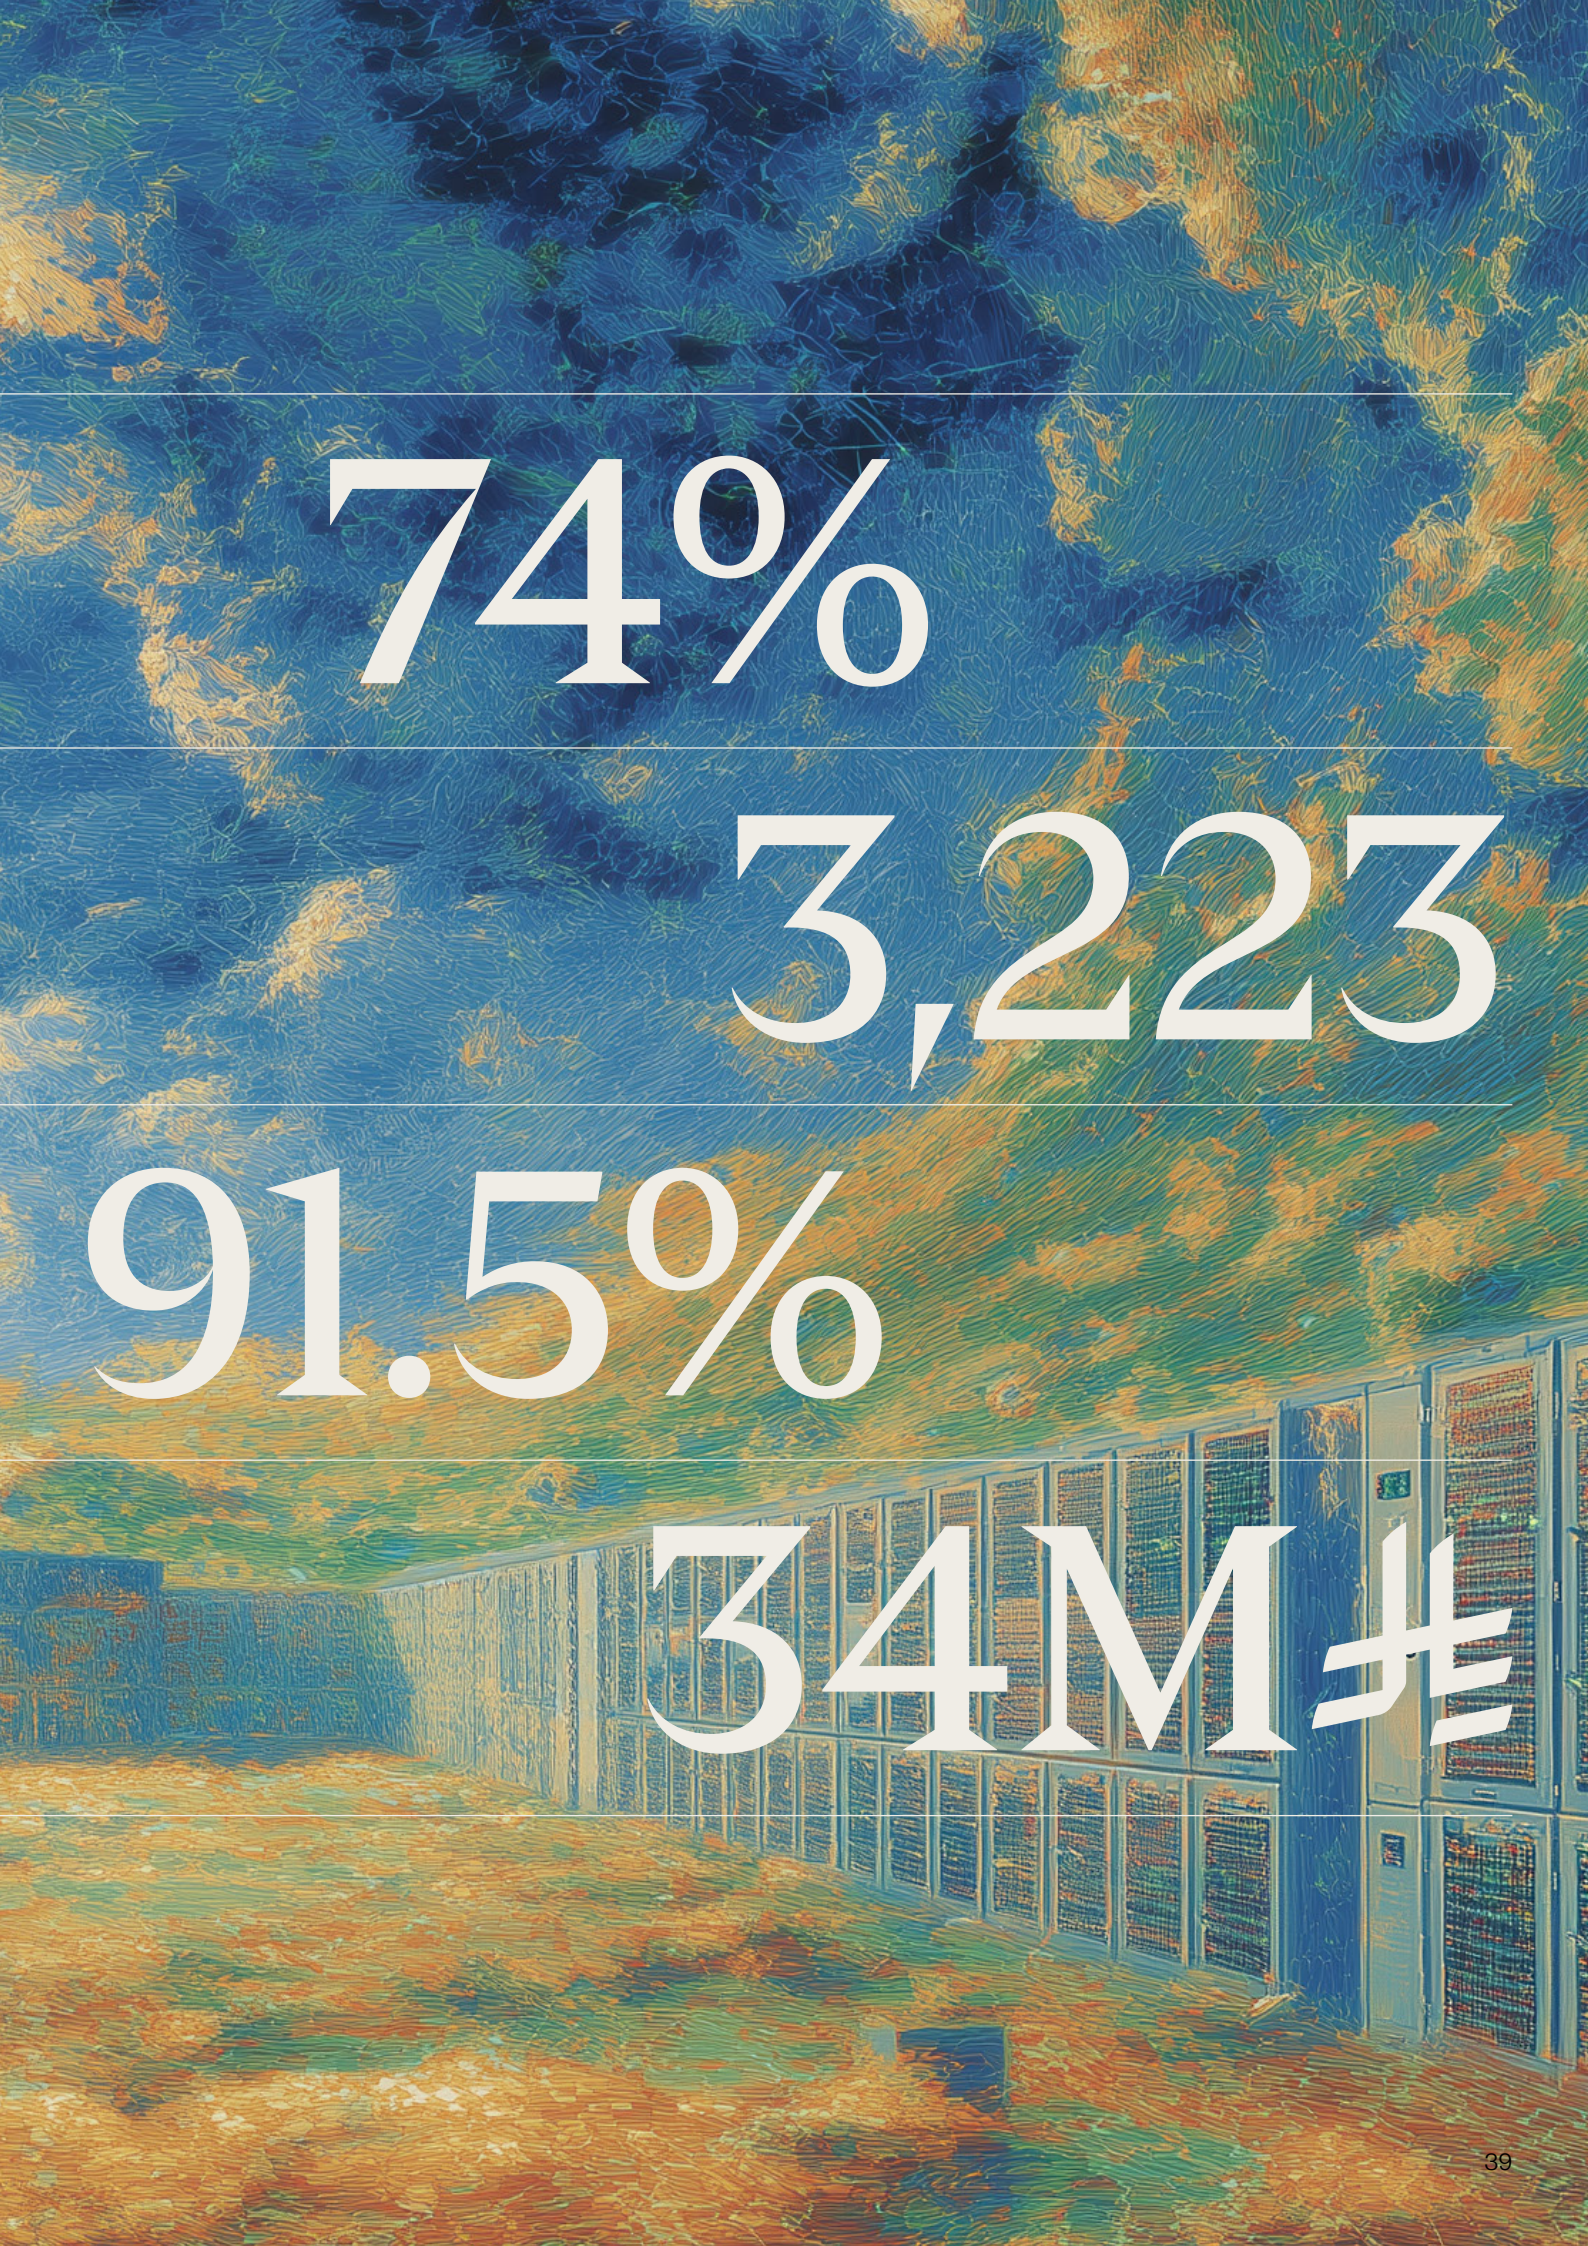
PEAK CONCURRENT CALLS - SINGLE CLIENT, SINGLE DAY

AUTONOMY RATE - FROM DAY ONE

SAVED ACROSS CLIENTS

Backed By

STV Google


74%
3,223
91.5%
34M≠

The Work

AI agents are only as valuable as the problems they solve. What follows is running in production today. Out of respect for our clients' operations, they go unnamed.

Embedded Everywhere

A national authority

One deployment we are proudest of A national authority brought Sawt inside its walls and put the agent across every channel it runs: the call line, the chat, the portal, the places its public actually shows up. One caller asks by voice, another types, a third leaves a message at midnight, and the same intelligence answers all three, with the same context and to the same standard.

This is what embedding means. The agent stops being a channel and becomes part of how the institution speaks. When a new service launches, the agent carries it everywhere at once, and the institution's people watch one view of every conversation the public is having with them.

Answering the Hospital's Phone

A government health network

Patients and visitors calling a hospital need answers. Appointment procedures, clinic hours, referral requirements, the small questions that decide whether a visit goes smoothly. Every one of those calls used to wait for a human agent, inside a system already carrying real pressure.

Sawt's agent now takes the inbound line across the network's hospitals. It answers at any volume, at any hour, with the same consistency on the thousandth call as on the first, and with the caller's context in front of it. The people it works alongside spend their time on the cases that need a person, which in a hospital is the point of everything.

Every operation has a conversation it cannot scale. Bring it to us.

The Thirty-Minute Window

A same-day gifting company

A gift usually goes to someone other than the buyer. That is the point of a gift, and it is also a problem: the recipient does not know a delivery is coming, and no one has told the driver the way to her door. Before anything moves, someone has to reach her and confirm the address, inside a narrow window, at volumes that surge exactly when everyone is celebrating.

Sawt's agent makes that call, with the order already loaded. It handles the hesitant recipient, the landmark instead of a street name, the wrong number, and hands to a human when a human is needed. The surge that once buried a team is now simply volume, and the gift arrives while it still means something.

The Hardest Call in Business

A consumer finance company

Debt collection is the most delicate conversation a company can have. The wrong tone at the wrong moment loses more than the payment. It loses the relationship.

Sawt's agent makes outbound collection calls with the account fully loaded before the first ring: the balance, the history, the terms. It opens with respect, holds the right tone through resistance, and knows when to press and when to leave room. Every call sounds like the company at its most careful, because that is exactly what it is.

Listening at Scale

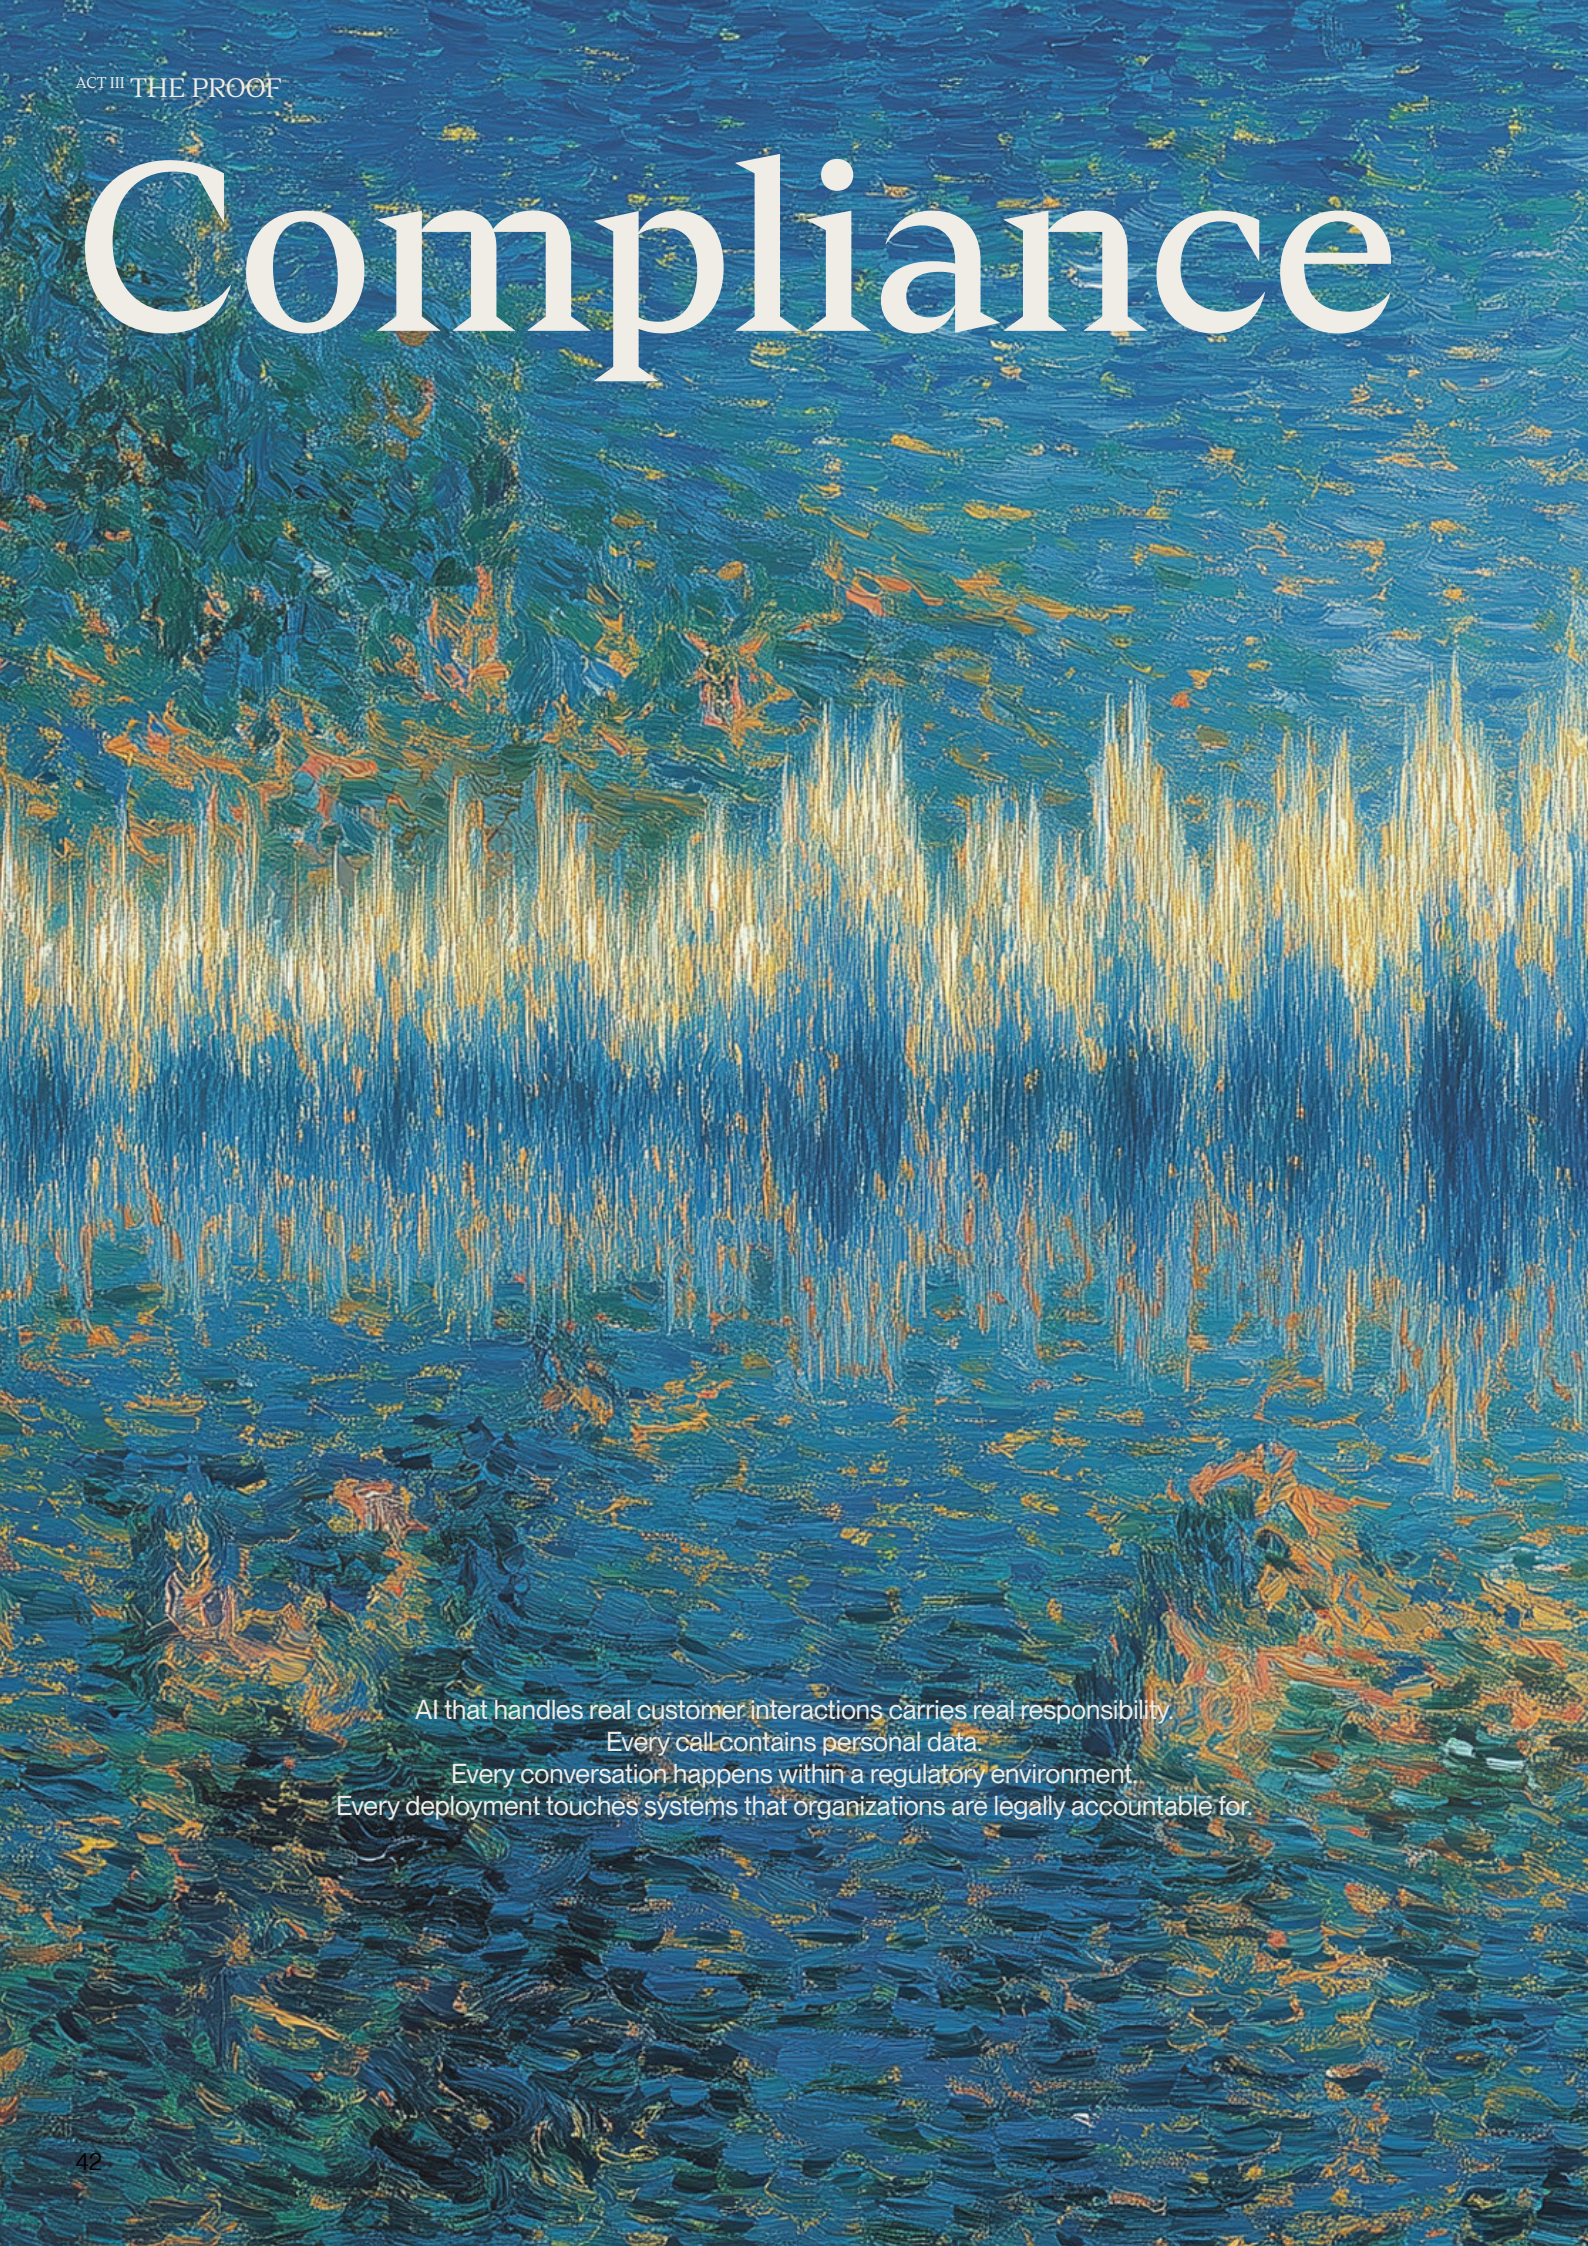
Research and outreach

Most surveys arrive as a form and die as one. Sent at the wrong time, ignored by most, answered by the few with the strongest opinions, they tell a company little about what its customers actually experienced.

Sawt's agent calls at the right moment, asks in a natural voice, and listens to the answers, including the ones that wander off the questionnaire.
The same discipline runs our lead qualification work: the agent takes the first conversation, learns what the person actually wants, and brings the sales team only the people ready to talk. In both cases the pattern holds. The repetitive conversation runs at scale, and the people move up a level.

Address confirmations, collections, surveys, qualification, a public line answered across every channel.
The use cases differ and the loop stays the same, because we sell the outcome of a conversation, whatever the conversation is.

Compliance



AI that handles real customer interactions carries real responsibility.
Every call contains personal data.
Every conversation happens within a regulatory environment.
Every deployment touches systems that organizations are legally accountable for.

Data Sovereignty

Sawt runs on-premise and is locally hosted within Saudi Arabia for organizations that require it. Your data does not leave your infrastructure — and it does not leave the Kingdom. No third-party cloud dependency, no cross-border transfer, no shared environment. You own the stack end to end.

PDPL

Saudi Arabia's Personal Data Protection Law governs how personal data is collected, processed, and stored. Sawt is built to operate within it — not around it.

ISO 27001

Our information security management is certified to international standards. The controls that protect your data are not internal policy — they are externally verified.

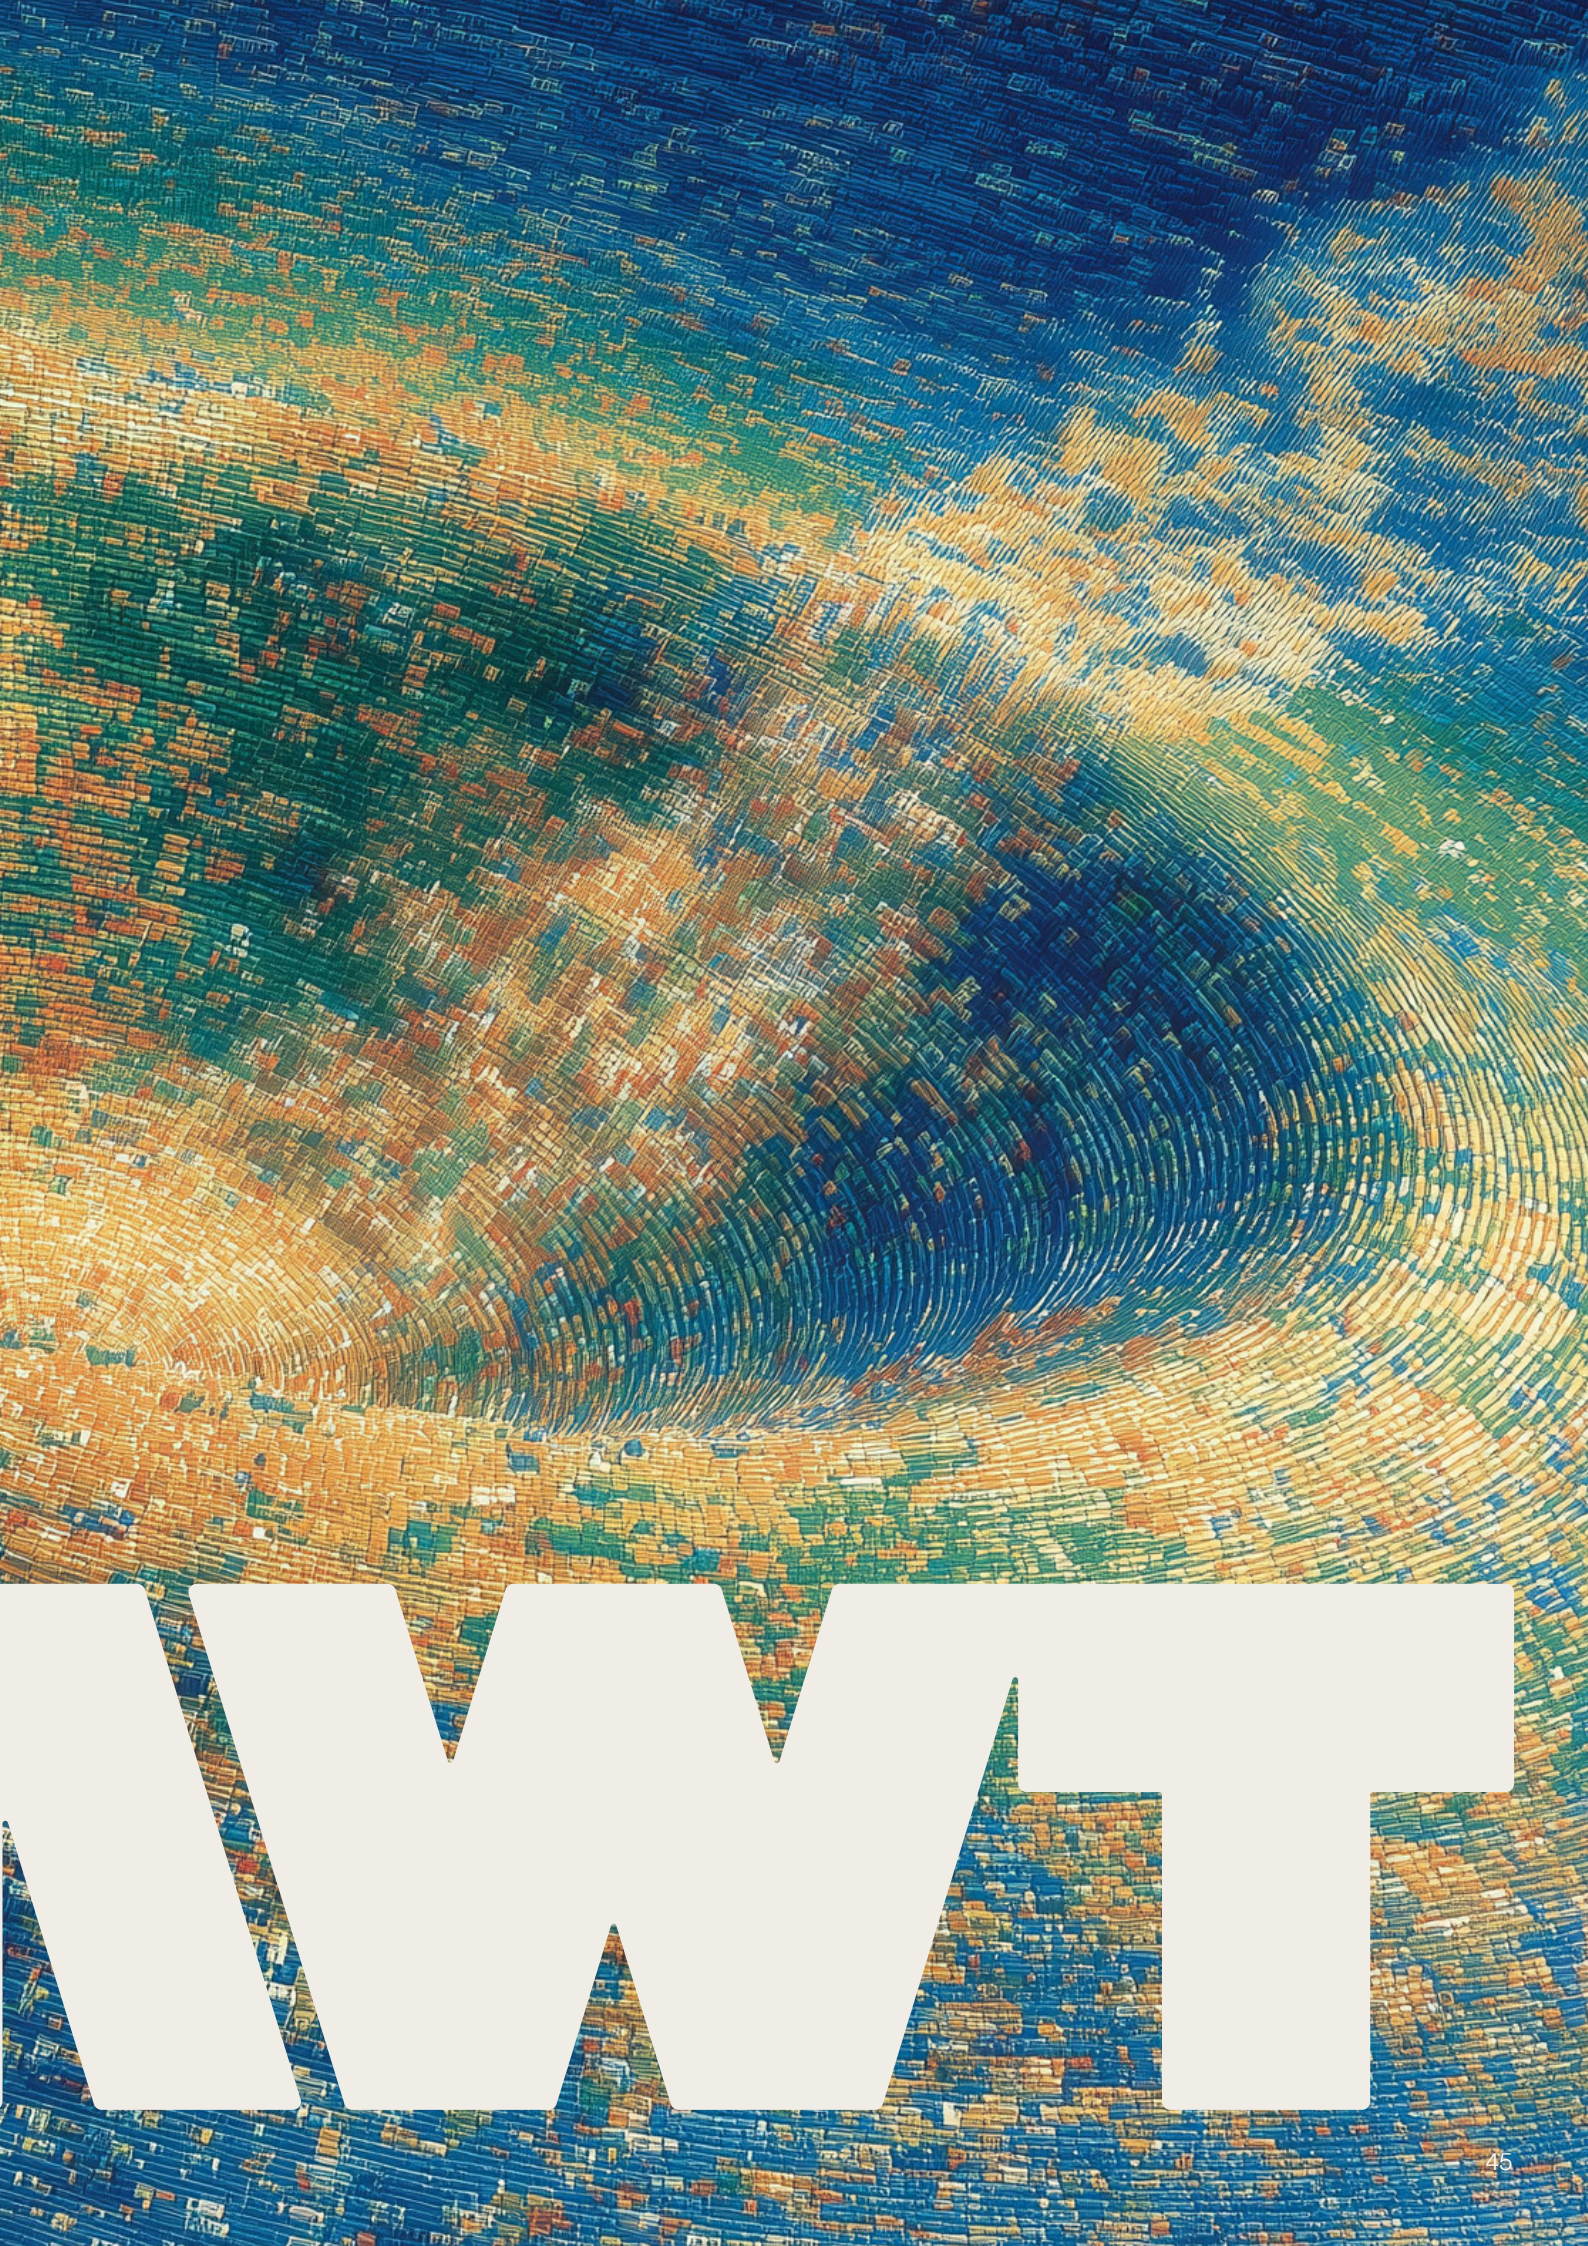
Built for Regulated Industries

Government. Healthcare. Finance. The sectors where a data breach is not just a reputational problem — it is a legal one. Sawt is deployed in all three.

Trust is not a feature you add at the end. It is the infrastructure everything else runs on.

ACT IV

Inside


WTF

Praxis

There is an old word for the moment an idea becomes work: praxis. Philosophers have argued about it for centuries. We use it for something simpler. What a company believes shows up in what its people spend their days doing, and nowhere else.

Here is what we believe. No one should be lost to meaningless work.

Every operation has work that numbs. The same question answered four hundred times before lunch. The dashboard watched all day for the one number that moves. The report assembled every Monday from the same four sources. Work like this asks nothing of the person doing it, and that is exactly the harm: work that asks nothing of a person slowly teaches them they have nothing to give.

We build software that takes that work. So the belief begins as a product. In every deployment, the agent absorbs the repetition and the people alongside it move up a level, to the cases that need judgment, patience, and taste. Ask the teams we work beside what changed and the honest answer is the shape of their day.

The belief also runs the company from the inside. An engineer here who meets the same task twice builds the tool that ends it. Nobody babysits a process that could run itself. The rule is the same at every desk: if the work is repetitive, it belongs to software. If it needs judgment, taste, or courage, it belongs to you.

This is not generosity. It is how quality happens. A person spending their whole mind on a hard problem produces work no divided person can match, and hard problems are the only kind we want to be known for. The craftsmanship this book keeps returning to has a precondition, and the precondition is attention, and attention is exactly what meaningless work destroys.

Machines should never get the interesting work.
People should never get the meaningless kind.

If the work is repetitive, it belongs to software. If it needs judgment, taste, or courage, it belongs to you.

The Team & The Obsession

Sawt is small on purpose. We hire for density, and we measure it honestly: every person here would be hard to replace, and knows the standard that got them in the door.

The people came back for this. Engineers and operators who sharpened themselves in the world's best markets and labs, and then chose Riyadh, because nowhere else does the work matter this much to the place it serves. They brought the habits of those markets with them, and the habits are the obsession this page is named for.

Nothing ships on a feeling. Every claim gets a measurement, every feature gets a bar, and production is the judge of both. When something fails, the failure is studied in the open and becomes the next test. The details a customer touches are treated as the product itself, because to the customer, they are. The team argues hard and holds opinions loosely, and the argument ends when the evidence arrives.

None of this is comfortable, and that is the design. Comfort and craft have never lived together long. What we offer instead is the thing craftspeople have always actually wanted: hard problems, real stakes, colleagues at their level, and work they would sign their name to.

The manifesto at the front of this book makes a promise about sophistication. Pages like this one are where the promise is kept. The standard our products carry has to be the standard our people set.

The Eight Values

01 Technical, always

Every problem, operational or not, gets attacked as an engineering problem.

02 Toe-stepping

No deference to hierarchy. If your view makes sense to you, fight for it.

03 Talent density above all

The roster is the strategy.
The best people sharpen each other.

04 Taste is the standard

Nothing ships, document, deck, or feature,
below the standard you would sign your name to. With AI, there is no excuse.

05 Tool literacy is table stakes

Up to date, and uses AI well.
In every role.

06 Ownership, start to finish

If people are chasing you for updates,
you are not communicating.

07 Written or it didn't happen

Decisions live where the whole team can read them.

08 High intensity, flexible in form

Outcome-obsessed.
The hours are yours; the outcomes are not.

Real AI

AI has made its way into almost every category of software. It writes, summarizes, answers, searches and automates. Much of it is genuinely impressive. But as AI becomes easier to build into a product, its presence tells us less and less about the value of that product.

A good demonstration can make intelligence look like the outcome. The model understands the question, responds naturally and completes the task exactly as expected. Production is less polite. Customers change direction halfway through a conversation. Policies carry exceptions. Internal systems take too long to respond. Demand arrives all at once. The same request comes phrased a hundred different ways. Intelligence has to survive all of it.

And the business still needs its outcome. The payment has to be collected. The appointment has to be booked. The issue has to be resolved. An operation that handled ten thousand interactions yesterday needs to handle twenty thousand tomorrow at something close to yesterday's cost.

This is the distance between using AI and building a business around what it can do.

The models will keep getting better, and access to them will keep getting easier. Intelligence that feels remarkable today will be available to everyone soon. When that happens, having AI will mean as little as having software means now.

The future of AI belongs to outcomes. The difficult work is moving past what a model can do in isolation, toward systems that hold up in production, improve through use, and can be trusted with a measurable part of the business.

Eventually we may stop describing products as AI at all. The technology will be assumed, and what remains will be what businesses have always cared about: what became faster, what became better, what became possible, and what it returned.

The most valuable AI will be the AI you no longer need to talk about.



Photography
Moath aloufi





The Customer Interaction Company.









Photography
Dana Awartani