

The **AI-Native** Organization

AI agents are software that does routine office work on its own: pulling figures out of a stack of PDFs, or drafting the comparison somebody writes by hand every month.

Competitors keep announcing what AI did for them. You want that here, before they get further ahead. **You cannot tell what is real.** This program builds that judgement, starting with you.

We run close to a hundred agents in our own operations. We have already made the expensive mistakes, including deleting the entire fleet once and starting clean.

START SMALL

Strategy consultation

USD 1,000

Ninety minutes on your own plan

BRING THE TEAM

The program

USD 8,000

Both working sessions, with your top team

FORMAT

Two working sessions,
90 minutes
online or in person

BETWEEN THEM

1–2 weeks of our
support

PARTICIPANTS

Your top team
priced per company,
not per seat

■ 01 · WHO YOU WOULD WORK WITH

The team

Everything below is checkable before you speak to us. We run close to a hundred agents inside our own operations, and we have already deleted the whole fleet once, because it turned into a mess, and started clean. Part of the method is filed: we have technology patent pending for adapting agents to the way a particular team works.

We measure the business, not the activity. More tickets closed, more code written, more drafts produced: none of that is a result. It counts when it shows up in something the company already counts, and we name that number before the work starts rather than after.

~100

agents running inside our own operations today

1

complete fleet deleted after it turned into a mess, then rebuilt clean

dozens

of setups tested on ourselves; a handful survived into daily use

**Dima Istomin**

Built Examus and sold it to Constructor Tech. Runs 8Hats day to day, with agents doing most of the work.

linkedin.com/in/dmitriy-istomin

**Mariam Mamedli, PhD**

Fourteen years in data intelligence for banks and fintechs. Head of Product, Asia at CEIC (ISI Markets), quants product line globally. PhD in Economics.

linkedin.com/in/mariamamedli

■ 02 · ONE ENGAGEMENT, TOLD STRAIGHT

A tech company of roughly 80 people: six months of systematic AI adoption, with subscriptions, pilots and an in-house ML team, and no visible effect on the business. The problem was not the technology. Top managers demanded adoption but did not work with agents themselves, so nobody could tell a good result from a plausible one. We started by onboarding the leaders. Afterwards one of them was able to judge his own AI work, and acted on it. Full case under NDA.

It is faster for the person who runs the business to learn what AI does than for a technical person to learn the business. That is why we start with the people who sign.

One engagement, client identity withheld. Not a guaranteed outcome.

Take a real announcement

A REAL ANNOUNCEMENT — KLARNA

“Our AI assistant handles two-thirds of customer service chats... doing the equivalent work of 700 full-time agents.”

THEIR WORDS, ON THE RECORD: [KLARNA.COM/INTERNATIONAL/PRESS](https://klarna.com/international/press). FROM THE OUTSIDE, YOU STILL CANNOT TELL WHICH READING APPLIES.

A sentence like that tells you nothing by itself. There are at least three readings.

READING 1

The number is true and empty

The chats it took were the scripted ones. Cost to serve did not move. The customer noticed nothing.

SAFE TO IGNORE

READING 2

Real work, wrong constraint

Genuinely deployed and genuinely working. Support was never what limited the business.

THEIR WIN. NOT YOUR MOVE

READING 3

Their economics moved

Cost to serve fell far enough that they can price or staff differently. This one you cannot ignore.

THE ONE TO ACT ON

Only the third should change your plan. Telling them apart is the judgement the program builds, and it applies twice: to what competitors publish and to what your own team reports back.

The first skill is telling a real result from a good story

It applies twice. To what you read about other companies, and to what your own team reports back to you. You cannot judge either from the outside. It takes having worked with agents yourself, on your own tasks, long enough to know what good looks like and what it actually costs to get. Until then every claim sounds equally plausible, and budgets get approved on tone of voice.

This is not a character flaw. Your people have automated real work already, and that is exactly the point. Automating a task is one problem. Running a company where people and agents share the work is a different one, and it is new enough that there is no playbook you failed to read.

■ 03 · WHAT WE RUN OURSELVES

Our own operation is what we put on the screen

About a hundred agents run inside our own firm today. In the working sessions that work is open on the screen, live, rather than described in slides.

Marketing

First draft of a page or a deck in an hour or two, not days and five rounds.

Research and analytics

Market research and buyer panels run overnight, before anything is built.

Sales support

Minutes and tasks from every call, ready as it ends. Any deal's history is one question away.

Development

Our own site and internal tools are written by agents; a person reviews before anything ships.

Correspondence

Agents write the drafts and turn calls into tasks; we read in the morning.

Money owed to us

Agents chase the invoices that are still open.

Checking that work takes minutes in the morning, and you watch us do it.



Our technology map at 8hats.ai/main#map lists what we use, with a label on each line: deployed and measured, working prototype, specified, or still frontier. Most suppliers show you the first category and let you assume the rest.


These are our own numbers, not a promise of yours. What a firm like yours can take from this is decided on your own material, in the two working sessions, and measured on a number your firm already keeps.

Where your client data goes. We work on material you choose and redact yourself. You see what each tool sends out, where it is stored, and the setting that stops your data being used for training, from the suppliers' own written policies. Nothing is connected to your systems for the working sessions.

■ 04 · WHAT HAPPENS

Three steps over three weeks

Before it starts, a short set-up call: we learn your context and collect the questions your people actually have. What comes out of it: a shared language, a working setup, and a review of what your team actually did with agents.



By "agent" we do not necessarily mean an autonomous agent. Any environment where a person works together with AI, Claude Code or Codex, already is a personal agent in our view.

■ 05 · WHAT YOU HAVE WHEN IT ENDS

What you have when it ends

- ✓ **A shared language and a decision frame.** What to delegate to agents, where a person stays, how to measure the result, understood the same way by the whole team.
- ✓ **Priorities for the first workflows.** One or two candidate processes, an owner and a metric. This is the basis for the pilot.
- ✓ **A setup for your team.** A concrete configuration: environment, model, skills, chosen from what we learn about your company rather than from a catalogue.
- ✓ **Materials from both working sessions,** the agreed priorities and the final setup, in writing.

■ 06 · WHAT CHANGES FOR YOU

The difference is what you can judge

Not a new tool stack. A leadership team that can tell a real result from a plausible one, in its own company and in everyone else's.

BEFORE

- Every AI claim, inside or outside, sounds equally plausible
- Spending is approved on enthusiasm and tone of voice
- “Are we behind?” has no answer anyone can defend
- The plan is discussed in the language of tools

AFTER

- ✓ You can name what a good result would look like before it is claimed
- ✓ Each initiative is judged against the plan, not against the pitch
- ✓ You know which of your bets moved and which held
- ✓ The team argues about mandates and metrics instead of tools

The part that stays with you

Whatever you decide about AI, the responsibility stays inside your firm. A supplier's contract usually caps what he owes you at what you have paid him over the last year. Advice you give a client stays with the person who gave it. No contract moves that to a vendor, and no consultant can sign it for you.

That is why the people who sign are the ones who should learn this, not only the people who install it. It is also why waiting is not neutral: these tools are already on your staff's phones, and if nothing is written down about what may go into them, that is your firm's current rule on AI.

■ 07 · PRICE

What it costs

USD 8,000

USD 8,000 per firm, however many of you are in the room. Software subscriptions your team keeps afterwards are not included.

Two working sessions of 90 minutes each, online or in person · your top team · preparation, adaptation of examples and materials included. Priced per company, not per seat, so bringing the whole leadership team costs the same as bringing three people. Quoted in US dollars. Date and payment terms are fixed in the invoice.

■ 08 · WHAT ELSE WE DO

Everything below is separate from the two working sessions

Nothing here is bundled in. The first three you can add at any point; the last three become possible once a first workflow exists and there is something worth measuring.

WORK	WHAT IT IS
Strategy consultation USD 1,000	Ninety minutes, one to one. What AI can and cannot take over in your firm, and what that is worth in hours and money. Written up afterwards, private.
AI-Native Pilot	The workflow chosen during the program, taken to a working state. Includes setting the whole thing up for everyone who took part.
Benchmark your agents	Agent reliability assessed against the same evidence-classed framework, on measured behaviour rather than a vendor's claim.
Stand up a test rig	A harness that runs your agents against your own cases repeatedly, so a change can be shown to help rather than argued about.
Build the agents	Where a workflow is worth automating and nobody internally has the time, we build it with your team rather than for them, so it stays yours.

Pilots start at USD 35,000 and are quoted before any work begins. **To start: a 20-minute call, details on the next page.**

Start with a 20-minute call

Twenty minutes, no charge. About a hundred agents run inside our own operations today, and we show you that work on the screen. We go through what you have seen others claim and what you have tried, then tell you where you actually stand and the first move you can defend. If we are not the right help, you get the reason and a pointer to what we would do instead.

calendar.app.google/2Nn9ui6bcnK7xWyz5

Prefer to write first? WhatsApp +65 8053 6400.

The pages that follow are reference. They are not part of the two working sessions. They show the frame we use, where a first workflow is usually chosen, and what a finished one looks like.

■ APPENDIX A · THE DIFFERENCE

The company that gets this right does not look like yours with agents added

It means the rules are written somewhere an executor can read them, whether that executor is a person or an agent. A person is asked only when no rule covers the case.

AXIS	TRADITIONAL COMPANY	WITH AGENT-ASSISTANTS	RUN ON AGENTS
Who works	people use software	people delegate individual tasks	agents carry the work; people own mandates
Working hours	office hours	people + AI assistants	people and agents on both sides, around the clock
Decision latency	hours to weeks	minutes to days	seconds to minutes when the decision is reversible
Organizing unit	function, role, queue	team plus a set of agents	mandate: goal, bounds, resources

The CEO should not make more decisions per minute. The CEO should design which decisions never need to reach them. If you cannot name a decision that should never reach you, there is no system yet. There is just you, faster and more tired.

■ APPENDIX B · A POSSIBLE CONTINUATION

Choosing the first workflows

The transition does not start everywhere at once. It starts with one or two processes, chosen on impact, data readiness, risk, and how fast the result can be verified.

Impact How much of the business the workflow actually touches	Data readiness Whether the inputs exist in a form an agent can read	Risk What breaks if it is wrong, and whether that is reversible	Speed of proof How fast you can tell whether it worked
Sales From the customer conversation to the quote and follow-up. Lots of repetitive document assembly, a high price for stale information, and a result that is easy to measure.	Product From feedback and data to verified decisions and requirements. The agent keeps the link between what was decided and what is actually in the materials and plans.	Operations Recurring reconciliations, reports and approvals, wherever information diverges between departments and manual hand-offs eat days.	

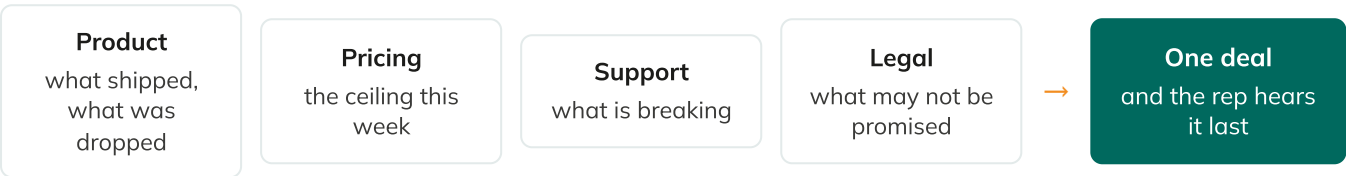
- ✓ **One workflow, one owner, one metric.** The pilot is scoped so the result can be proven, not just felt.
- ✓ **Observe first, act second.** Start in shadow or read-only mode; autonomy rises only after quality has been measured.
- ✓ **Narrow scope, complete cycle.** Half a cycle is worth nothing: if rulings do not reach the executors, and the exceptions never come back to change the rules, you have bought a document rather than a control system.

Why sales first. Sales is the only workflow where improvement has unbounded upside, because it grows revenue rather than cutting cost, and savings have a ceiling while sales do not. One mandate in sales immediately touches marketing, product, support and sales ops, so the whole operating model gets tested on it.

■ APPENDIX C · A WORKED EXAMPLE

What the sales workflow looks like afterwards

Deals are dropped because the information the rep depends on arrives last: the signals a deal depends on are born in four different departments and reach the rep last. We gather those signals into one workflow and turn them into precise instructions for every salesperson.



TODAY	WITH THE WORKFLOW IN PLACE
The morning starts with digging through email and guessing what changed	A morning feed: three things changed for your accounts, the promo runs to the 15th, plan X was replaced, the discount ceiling is now 12%. Thirty seconds to read.
A discount question waited a day for the manager's answer	Answered in seconds from the current rule, without asking anyone.
A quote goes out with last year's price or a discontinued module	Every quote is checked against the model before it is sent, with discrepancies highlighted.
A promise lives longer than the product it was built on	A risky proposal is held and routed to a person. This is the case that should reach a human.

What we measure. An improvement only counts if the metric was named before and measured after: time from conversation to a ready quote, share of quotes approved on the first pass, and the number of discrepancies caught before the customer sees them.