

ieYou Tech

亦友科技 · ENTERPRISE AI

COMPANY PROFILE

Your business, running on **AI**.

We don't sell tools, and we don't sell slide decks. We turn your data, your workflows and your rules into a layer AI can actually operate on — then build the systems on top of it, and stay until your own team runs them without us.

ieYou Technology (Ma'anshan) Co., Ltd.

亦友科技(马鞍山)有限公司 · Established June 2023
Offices in China and Japan

2026 EDITION

What's in this document.

Company at a glance Who we are, where we operate, what we deliver	01
Why enterprise AI stalls The five places AI programmes actually break down	02
Our method: a foundation first Why the layer underneath AI decides the outcome	03
Five service lines Strategy · Embedded engineering · Applications · Private deployment · Data governance	04
Industry depth: scientific research and testing Data governance where it is hardest	05
Sovereign AI Architecture patterns for on-premise deployment	06
Selected engagements Five anonymised projects delivered by ieYou	07
Experience our team brought with them Transformation and early generative AI work led before ieYou	08
Where we operate China and Japan	09
How we work, and how we charge Engagement flow and commercial models	10
Leadership Who you will be working with	11

A note on confidentiality

Every engagement in this document has been anonymised. Client names and commercially sensitive details are withheld or generalised in line with the confidentiality agreements under which the work was performed. Where a project was led by a member of our team before ieYou was founded, we say so explicitly rather than presenting it as company work.

The engagements shown are a selection, not a complete record, and are not the limit of the industries or problem types we work in. Further experience can be discussed under NDA.

An AI delivery firm, not an advisory practice.

Established in June 2023, we help mid-sized enterprises and research institutions get AI into daily operation — starting from the data and process foundation that almost every failed AI programme skipped.

June 2023

Established as ieYou Technology (Ma'anshan) Co., Ltd. — 亦友科技(马鞍山)有限公司

China & Japan

Headquarters in Ma'anshan, Anhui. Japan operations with a local partner under the ZENKAI OS brand.

Five service lines

From deciding what is worth building, through to running systems your team owns.

What we actually do

Most organisations approaching AI have already tried something — a pilot that impressed everyone in the demo and then quietly stopped being used. The gap is rarely the model. It is that the business's data, processes and rules were never put into a form software could operate on reliably.

That is our work. We map and standardise the underlying data, redesign the process where it needs it, build the systems on top, and — where data cannot leave the building — deploy the models on the client's own hardware. The engagement ends when the client's team can run the result without us.

Who we work with

- **Mid-sized enterprises** that know they need AI but cannot get past the pilot stage.
- **Research and testing institutions** with fragmented instrument and experiment data.
- **Organisations under strict data rules** for whom public-cloud AI was ruled out on day one.

Corporate details	
Legal entity	ieYou Technology (Ma'anshan) Co., Ltd. · 亦友科技(马鞍山)有限公司 · established 1 June 2023
Headquarters	Ma'anshan, Anhui Province, China (full address on page 16)
Japan	Shizuoka City · operating with a local partner under the ZENKAI OS brand
Team	A small senior team of engineers and consultants. No junior bench, no subcontracting of core work.
Contact	chase_dai@ieyoutech.com · ieyoutech.com

Enterprise AI rarely fails on AI. It fails in five predictable places.

We have yet to meet a client whose AI difficulties were caused by choosing the wrong model. The pattern is consistent enough to list.

- 01 The strategy arrived, and ownership did not.** A capable plan exists on paper. Nobody can say who does the next thing, by when, with what budget.
- 02 The pilot demonstrated beautifully, then met real work.** Curated data and a rehearsed scenario hide every edge case that makes the process expensive in the first place.
- 03 The system went live, and the team went back to spreadsheets.** Adoption was treated as training rather than as process design. The new tool asked for more effort than the old habit.
- 04 The data disagrees with itself.** The same customer, product or sample exists in several systems under several definitions. No model can reconcile what the organisation has never reconciled.
- 05 The data cannot leave the building.** For regulated institutions, public-cloud AI is not a preference to be debated. It is excluded in the first meeting, which ends most vendor conversations.

Why we organise our services this way

Our five service lines map onto these five failures deliberately. Strategy work exists to produce ownership, not documents. Embedded engineering exists because plans written by people who never build them do not survive contact with real data. Applications are built against one specific process rather than as general assistants. Private deployment exists because for many clients it is the only route that is legally available. And data governance comes first, because everything above it depends on it.

What this means commercially

Clients who have already been through a failed pilot are usually not looking for more ambition. They are looking for something narrow that demonstrably works, delivered by people who will still be there when it breaks. That is how we scope first engagements: one problem, a few weeks, something real to look at, and an honest answer about what should come next — including when the honest answer is that AI is not the right tool for the problem in front of us.

A foundation first. Then AI.

An AI system is only as capable as the layer beneath it. Without a reliable representation of the business, a language model is a well-spoken interface with no authority to act.

What we mean by the foundation

The foundation is the part of the organisation that software can read and act on: entities defined once and consistently, processes described precisely enough to be executed, rules made explicit rather than held in the heads of experienced staff, and a record of what happened that can be trusted afterwards. Building it is unglamorous and it is the difference between AI that answers questions and AI that completes work.

01 • A foundation first, then AI

We turn data, processes and rules into a layer AI can operate on. Without it, AI is a chat window. With it, AI can carry real work — place an order, close a case, flag an anomaly, produce a report that someone acts on.

02 • Engineers on site, no subcontracting

The people who write the plan are the people who build it. They work from the client's office, with the client's data, and hand over to the client's team rather than to another vendor.

03 • Built at a size you can afford

This way of working has historically been reserved for governments and the largest enterprises. We deliver it at a scale a mid-sized organisation can actually commission: a single problem, a few weeks, a working result.

04 • An honest scope

We would rather remove a use case from scope than deliver something nobody uses. Part of the first engagement is establishing which problems AI should not be pointed at yet.

Our position, stated plainly

We are not a software vendor and we have no platform product to sell you. What we build belongs to the client, runs in the client's environment, and can be maintained by the client's staff. Where we reuse our own technical components — for example our attestation infrastructure — they are deployed into the client's estate, not rented back to them.

01 AI Strategy

Establishing what is genuinely worth building

We look at operations before we discuss technology. Where time is lost, where errors are introduced, where margin disappears — and which of those AI can address now, later, or not at all. The output is a plan that can be started, not a report to be filed: what to build, who owns it, what it costs, and when it pays back.

- **Operational review** — process mapping, quantified pain points, current-state data assessment.
- **Use-case ranking** — feasibility against value, sequenced rather than listed.
- **Payback modelling** — investment, expected return, and the assumptions each depends on.
- **Roadmap and ownership** — named owners, decision points, and what must be true before each phase starts.

2–6 weeks typical

Board-ready output

Can be delivered standalone

02 Embedded Engineering

Forward deployed engineers inside your team

An engineer joins the client's team, works with their real data and real process, and ships continuously. Weekly or monthly cadence, on site where the work requires it. The engagement is complete when people are genuinely using what was built — not when the code is handed over.

- **On-site delivery** — the engineer sits with the people whose work is changing.
- **System integration** — connecting to the ERP, CRM, LIMS or line-of-business systems already in place.
- **Process redesign** — changing the workflow where the workflow is the problem.
- **Capability transfer** — documentation, pairing and handover as part of delivery, not as a final phase.

Monthly or quarterly

Same team throughout

No subcontracting

03 AI Applications

Systems people use every day

We do not build general-purpose chatbots. Each application targets one specific process: search and question answering across the organisation's own documents, contract and report processing, automation inside approval flows, or multi-step agent workflows that carry a task from start to finish. A usable version ships first; iteration follows from real use.

- **Document AI** — extraction, classification and processing of contracts, reports and forms.
- **Knowledge search** — retrieval over internal material, with answers traceable to source.
- **Process automation** — AI placed inside an existing approval or handling flow rather than beside it.
- **Agent workflows** — multi-step tasks with tool access, integrated into the systems of record.

First release in weeks

Built against one process

Adoption measured, not assumed

04 Private Deployment & AI Infrastructure

AI that runs inside your own building

Open-weight models deployed on the client's own servers. We size and procure the hardware, build the environment, and operate it through the first phase of production. Nothing leaves the client's network. For organisations under hard data rules this is usually the only route that is available at all — and increasingly it is chosen for control rather than compliance.

- **Model selection and deployment** — current open-weight model families, chosen against the actual task rather than benchmark position.
- **Hardware sizing and procurement** — GPU specification against real data volume and concurrency, purchased and configured by us.
- **Inference operations** — serving, monitoring, capacity and upgrade path.
- **Network segregation** — architectures in which collection and analysis are separated, with no traffic transiting the public internet.

Fully on-premise

Hardware procurement included

Air-gapped patterns available

05 Data Governance & Trusted Data

Getting data to a state AI can use — and prove

First the inventory: what data exists, who owns it, in what format, how it moves. Then standards, pipelines and quality gates, so that AI has something reliable to work from. Where the provenance of a record matters as much as its content, we add attestation, so that what happened to the data can be demonstrated later rather than asserted.

Governance

- **Data inventory** — sources, owners, formats, flows.
- **Standards** — master data definitions that the people entering data will actually follow.
- **Pipelines** — collection and integration that runs unattended.
- **Quality control** — validation at the point of entry, not as a later clean-up.

Trusted data

- **Attestation** — records and metadata committed to a permissioned chain at the moment of creation.
- **Provenance** — a verifiable history of who did what to which record, and when.
- **Tamper evidence** — alteration after the fact is detectable rather than merely discouraged.
- **Smart contracts** — rules enforced by the infrastructure instead of by policy documents.

Our attestation infrastructure

We operate our own permissioned chain infrastructure, built on open-source blockchain components and refined across multiple client deployments. It is a technical foundation we deploy into the client's environment — not a platform product or a subscription service. Two of the engagements described later in this document run on it.

We are deliberate about where this belongs. Attestation is valuable when provenance carries real consequence: scientific records that may be audited years later, operational logs in a regulated activity, evidence that must hold up outside the organisation that produced it. Where it adds cost without adding assurance, we say so.

Usually the first engagement

Prerequisite for everything above it

Deployed, not licensed

Scientific research and testing: data governance where it is hardest.

Laboratory data is among the most difficult material to govern that exists in commercial practice. It is also where we built our methodology.

Why laboratories defeat generic frameworks

Instruments export in a dozen incompatible formats, several of them undocumented. Every technician writes up results in their own way, and each way is defensible. Records live simultaneously in a LIMS, in spreadsheets on individual machines, and on paper. Projects are accounted for by funding line rather than by process, so the same physical experiment appears under different identities depending on who is asking. A governance framework designed for customer or financial data does not survive contact with this environment.

What we bring

- **No ramp-up on the domain.** Instrument interfaces, laboratory workflow and testing standards are familiar to us. We do not spend the first month learning what a sample is.
- **Governance aimed at AI, not at the shelf.** The objective is a foundation that AI systems can rely on — not a policy binder that satisfies an audit and changes nothing.
- **Project accounting that matches reality.** We model projects, samples, instruments and personnel the way the institution actually organises them, which is what makes adoption possible.
- **Provenance built in.** In research, a result that cannot be traced is a result that cannot be defended. Attestation at the point of capture makes traceability a property of the system rather than a reporting exercise.
- **We hand over running pipelines.** Working code and operating procedures, not documentation describing what should be built.

Beyond research and testing

Research and testing is where our data methodology was developed and where our depth is greatest, but the underlying work — reconciling fragmented sources, defining standards that people will follow, and building pipelines that run unattended — applies wherever operational data is scattered across systems that disagree. Our engagements to date also include consumer goods, low-altitude aviation operations, and digital content production.

Keeping intelligence inside your walls.

A growing share of the organisations we speak to cannot send their data to a public AI service — and a growing share of those who legally could have decided they would rather not.

What on-premise AI means in practice

It means open-weight models running on hardware the organisation owns, in a network the organisation controls, with no inference traffic leaving the perimeter. It also means someone has to make unglamorous decisions correctly: how much GPU memory the workload genuinely needs, how the model is served, what happens when it has to be upgraded, and how the sensitive and non-sensitive halves of a system communicate.

A pattern we deploy: split collection and analysis

Where data must be gathered from external or internet-facing sources but analysed in confidence, we separate the two concerns physically. Collection services run on internet-facing nodes and do nothing but acquire and forward. Transfer to the secure environment is point-to-point, without transiting a third-party service. All analysis, summarisation and generation happens on locally deployed models inside the protected network. The externally exposed component holds no analytical capability and no accumulated intelligence; compromising it yields collection plumbing and nothing else.

What we take responsibility for

- **Sizing.** GPU specification against real data volume and concurrency, rather than against a vendor's recommendation.
- **Procurement.** We buy and configure the hardware. Clients without infrastructure teams do not have to become buyers of GPUs.
- **Model selection.** Current open-weight families, chosen against the specific task. Where a smaller model is sufficient, we use the smaller model.
- **Operations.** Serving, monitoring and the upgrade path, through the first phase of production and into handover.

We are candid about scale. Our deployments are sized for the workload in front of them, which for most mid-sized organisations means a modest, well-configured environment rather than a cluster. Being honest about that is part of why the numbers work for clients this size.

Anonymised project experience.

Delivered by ieYou Technology. Client identities withheld under confidentiality agreement. A selection, not a complete record.

TRUSTED DATA · AVIATION OPERATIONS

Verifiable operational control for a drone operations group

CHALLENGE

The operator needed its flight records to be defensible to third parties. Flight, mission and pilot data lived in separate operational systems, and records that could in principle be edited after the fact were of limited value as evidence.

WHAT WE DID

Synchronised flight data, mission data and pilot credentials into a single operational record, then committed that record to a permissioned chain through smart contracts at the point of capture. Attestation rules are enforced by the infrastructure rather than by procedure.

OUTCOME

Operational records became tamper-evident and independently verifiable. Flight history, mission authorisation and pilot qualification can be demonstrated for any given sortie, and because attestation happens at the point of capture the record is complete by construction rather than by later reconciliation.

DATA GOVERNANCE · SCIENTIFIC RESEARCH

Scientific data management platform for a research institution

CHALLENGE

Experimental data was produced by instruments in incompatible formats, recorded inconsistently by individual researchers, and dispersed across a LIMS, local spreadsheets and paper. The institution needed data that could be found, trusted and reused — and traced years later.

WHAT WE DID

Began with a survey of existing infrastructure and network conditions, then designed the experimental data collection scheme around it. Modelled projects and data in the institution's own project accounting terms, allowing researchers to associate specific instruments with each project and define per-instrument data templates. Experimental data now synchronises automatically to a distributed storage platform under central control, with metadata committed to the chain at the point of capture and every operation attested.

OUTCOME

A single governed record of experimental data with verifiable provenance for future audit and reuse. Researchers no longer transcribe results between systems: instrument output reaches governed storage without a manual step, and any operation on a record can be reconstructed afterwards. AI can now be applied to the institution's own scientific material rather than to a cleaned-up extract of it.

PRIVATE DEPLOYMENT · CONFIDENTIAL ANALYSIS

Multi-source data aggregation and analysis system, deployed on premise

CHALLENGE

A Beijing-based client required a system that aggregated data from many external sources and analysed it, under a confidentiality requirement that ruled out any commercial AI service and any analysis outside their own environment.

WHAT WE DID

Separated collection from analysis. Collection services run on internet-facing nodes; data is forwarded point-to-point into the protected environment, where locally deployed open-weight models perform analysis, summarisation and downstream processing. We specified and procured the hardware and built out the inference environment.

OUTCOME

Full analytical capability with no data or inference traffic leaving the client's control, and no analytical capability resident on the externally exposed nodes.

AI APPLICATIONS · DIGITAL CONTENT

Agent-driven creative production system for a content studio

CHALLENGE

A production team working across scripts, images and video was moving constantly between disconnected tools, with no shared workspace in which a project could be developed from concept through to finished assets.

WHAT WE DID

Built an infinite-canvas production environment integrating multiple generation APIs behind a single interface, with built-in production skills — character turnaround sheets, art-style transfer, camera movement and others — available as composable operations. A conversational agent sits inside the canvas, so that work can be directed through dialogue rather than assembled step by step.

OUTCOME

Script, image and video work happens in one workspace. Multi-step generation that previously meant moving between several disconnected tools is now directed through dialogue, and because production conventions are encoded as reusable skills, a decision about style or framing has to be made once rather than repeated on every asset.

CAPABILITY TRANSFER · JAPAN

AI adoption programme for our Japanese partner organisation

CHALLENGE

The partner's team was interested in AI but had no practical basis for applying it to their own work. Buying a tool would not have addressed the underlying issue, which was that nobody knew what to build or how to specify it.

WHAT WE DID

Delivered a structured AI application programme aimed at capability rather than familiarity: how to identify a workflow worth automating, how to build an assistant against it, and how to iterate once it is in use.

OUTCOME

The team now writes its own AI assistants without our involvement. Daily email is classified, summarised and answered by assistants they built themselves; scheduling and travel booking are handled the same way. New use cases are identified and built by the partner's own staff, without a request coming back to us.

WHY THIS MATTERS TO US

This is the outcome our third principle describes. The client did not become dependent on us — they became capable without us. It is the result we aim for in every engagement, and the one that is hardest to sell in advance.

What these engagements have in common

They look unrelated on the surface. What connects them is the shape of the problem: in each case the client's data or process had to be put into a usable state before anything intelligent could be built on it, and in each case the engagement was judged on whether the result was still in use afterwards.

Further engagements, including work we cannot describe even in anonymised form, can be discussed under NDA. The industries above reflect what we have been asked to do so far rather than a defined sector focus.

Work led before ieYou was founded.

The projects below were led by our founder at an international consulting firm prior to establishing ieYou. We describe them because they are the origin of how we work, and because clients are entitled to know where the method came from.

DIGITAL TRANSFORMATION · CONSUMER GOODS

Sales and trade marketing digitalisation for a national dairy group

CHALLENGE

Nationwide distribution and in-store execution depended on manual reporting. Master data was defined differently across systems, and trade marketing spend could not be attributed to outcomes with any confidence.

WHAT WAS DONE

Mapped the end-to-end sales process and identified where it broke, established master data standards across the affected systems, and designed a digital trade marketing model built on image recognition and mobile field applications. Coordinated the sales force automation and trade promotion management vendors, and produced the technical solution integrating the two systems.

OUTCOME

Manual reporting left the field workflow, and in-store execution became verifiable through image recognition rather than self-declared. With master data reconciled across both platforms, trade marketing spend could be attributed to specific activity for the first time — which is what made the cost reduction during the programme possible.

GENERATIVE AI · CONSUMER GOODS

Autonomous customer service agent for a global consumer goods group

CHALLENGE

Customer service handled a high volume of routine transactional contact across several channels and three languages, with agents spending most of their time on order status and logistics enquiries that the systems of record could already answer.

WHAT WAS DONE

Designed a fully automated agent handling live chat, email and voice telephony in Mandarin, Cantonese and English — and, critically, able to transact rather than only respond: placing orders, answering order status enquiries and tracking shipments through integration with the order management and CRM systems, updating order and customer state directly.

OUTCOME

Designed and approved for delivery. Our founder led the solution design and client engagement; implementation proceeded after his departure from the firm.

GENERATIVE AI · LUXURY RETAIL

Real-time virtual try-on for a European luxury fashion group

CHALLENGE

The group wanted customers to see new season pieces on themselves immediately, in store and in its mobile channel, and to convert browsing into styling advice rather than a catalogue experience.

WHAT WAS DONE

Designed a real-time garment visualisation capability combining large language and diffusion models, generating the customer wearing selected pieces, and recommending complete looks based on their purchase history and stated preferences — with the recommended look rendered rather than merely described.

OUTCOME

Designed and approved for delivery, targeting both in-store experience and conversion rate. As with the engagement on the previous page, implementation proceeded after his departure from the firm.

Why 2023 solutions are still relevant evidence

Both generative AI engagements in this section were designed in 2023, when almost no enterprise had put a language model into a transactional workflow. The tooling available then was considerably worse than what we work with today, which is precisely the point: the hard part was never the model. It was grounding an agent in the systems of record and giving it authority to act rather than only to answer. That is the same problem we solve for clients now, and it is why the sequencing described earlier in this document — foundation first, then AI — is not a stylistic preference.

Why we present this experience separately

It would be straightforward to fold these projects into the previous section and let them read as company work. We do not, for two reasons. They were delivered under another firm's contracts, and clients who later discover that are right to discount everything else in the document. And the value in them is not the logo — it is that the person who designed them is the person who will be in the room on your engagement.

China and Japan.

Two locations, both operational. We would rather describe two offices accurately than list ambitions as presence.

China · Headquarters

ieYou Technology (Ma'anshan) Co., Ltd. is headquartered in Ma'anshan, Anhui Province, and serves clients nationwide. Delivery is on site: our engineers travel to the client and work from the client's premises for the duration of the engagement.

Address	Room 1205, Building B, Innovation & Entrepreneurship Base, Southeast University Science Park, Huashan District, Ma'anshan, Anhui Province, China 安徽省马鞍山市花山区东南大学科技园双创基地 B 栋 1205
Contact	chase_dai@ieyoutech.com

Japan · Partner operations under the ZENKAI OS brand

In Japan we operate with a local partner under the ZENKAI OS brand, based in Shizuoka City. The partnership exists because AI adoption in Japan is a local business: it depends on relationships, on language, and on working within Japanese organisational practice rather than importing a delivery model from elsewhere.

Our contribution to the partnership so far has been capability building — the adoption programme described on page 13, after which the partner's own team began building and running AI assistants independently. Engineering capacity from China is available to the partnership for delivery work that requires it.

Address	Libro BLDG 301, 1-1-8 Nakada, Suruga-ku, Shizuoka City, Shizuoka Prefecture, Japan 静岡県静岡市駿河区中田1丁目1-8 リブロBLDG 301
Brand	ZENKAI OS
Contact	chase_dai@ieyoutech.com

Working with clients outside China and Japan

We work with organisations outside these two markets, and our delivery model travels reasonably well: engagements are scoped around on-site periods rather than continuous presence, and the private deployment work in particular suits clients whose data cannot leave their own jurisdiction. Where a long-running engagement requires local presence, we would establish it rather than claim it.

From the first conversation to your team running it alone.

STEP 01 Talk The first conversation is free. Tell us what is slowing the business down. We will tell you where AI helps and where it does not.	STEP 02 Plan Priorities, payback and technology choices. Something that can be started on Monday, not a report to be filed.	STEP 03 Build on site Engineers join your team and work with real data and real workflow until the system runs in production.	STEP 04 Hand over Documentation, training and operational handover. The engagement ends when your team runs it without us.
---	--	--	---

Commercial models

Model	How it works	Suits
Fixed-scope assessment	Defined deliverable and timeline. Operational review, data assessment or use-case prioritisation.	First engagements, and organisations that need a decision before a budget.
Embedded engineering retainer	Monthly or quarterly, one or more engineers working inside your team with an agreed cadence of delivery.	Programmes where scope will legitimately evolve as real data is encountered.
Project delivery	Fixed scope and price against an agreed specification, with acceptance criteria defined up front.	Well-understood builds: a specific application, a defined deployment.
Infrastructure build	Hardware specification, procurement and environment delivery, with an optional operations period.	On-premise AI deployments, particularly first ones.

How a first engagement usually looks

One problem, a few weeks, a working result, and an honest assessment of what should follow. We prefer to start narrow, because a client who has already lost money on a pilot is owed evidence before ambition. Larger programmes are easier to justify once something is demonstrably in use.

Who you will be working with.

ieYou is deliberately small and deliberately senior. The person who scopes your engagement is the person who will be in the room while it is built.

Chase Dai • Founder

Chase founded ieYou Technology in June 2023. His background is in digital and AI transformation for large enterprises: process and data work in consumer goods, and generative AI solution design in customer operations and retail experience. The three engagements described on pages 14 and 15 were led by him prior to founding ieYou.

The firm exists because of a gap he kept encountering: strategy work produced plans that nobody could execute, and engineering vendors built systems that did not fit the business. ieYou is structured so that the same people do both — which is also why it does not scale by subcontracting.

Team and delivery capacity

ieYou runs as a small senior team of engineers and consultants. There is no junior bench and no account layer between the client and the people doing the work. The constraint this creates is deliberate — it is what makes the embedded model possible at all.

A typical engagement is led by one senior owner with one or two engineers on site. We manage capacity by taking on fewer engagements at once rather than thinning the team on each one, and where a project needs a specialist we do not have, we bring one in under our own accountability rather than subcontracting the scope. In Japan, delivery capacity is shared with our partner organisation under the ZENKAI OS brand.

We will tell you in the first conversation if an engagement is larger than we can staff properly. A firm this size that accepts everything is a firm that under-delivers somewhere, and it is cheaper for both of us to establish that before a contract than after one.

How we staff an engagement

- **One senior owner throughout.** No handover from a sales team to a delivery team, because there are not two teams.
- **Engineers on site.** Delivery happens where the data and the process are, not remotely against a specification.
- **No subcontracting of core work.** Where we bring in specialists, they work as part of our team under our accountability.
- **Handover as a deliverable.** Capability transfer is scoped and priced into the engagement, not offered as an afterthought.

ieYou Tech

Thirty minutes to find out whether AI can actually work in your business.

You do not need a requirements document. Tell us what is not working, and we will tell you honestly whether this is a problem worth pointing AI at — including when it is not.

Email · chase_dai@ieyoutech.com

Web · ieyoutech.com

China · Ma'anshan, Anhui · Japan · Shizuoka City

ieYou Technology (Ma'anshan) Co., Ltd. · 亦友科技(马鞍山)有限公司

© 2026 ieYou Technology. All rights reserved. This document contains anonymised client information and is provided for evaluation purposes.