

The **AI-Native** Playbook for Hong Kong Businesses

How to choose, pilot and scale a first AI project on Microsoft Azure
— without the hype.

ICO BROTHERS LIMITED

Room 602, 6/F, Kai Yue Commercial Building, No. 2C Argyle Street, Mongkok, Kowloon, Hong Kong

Edition 1 · 2026

What is in this playbook

1. What "AI-native" means in practice
2. Why bolt-on AI features disappoint
3. Choosing a first use case: the scoring sheet
4. The four-to-six-week pilot
5. Data, privacy and PDPO questions to settle first
6. The Azure building blocks
7. Readiness checklist
8. About ICO BROTHERS LIMITED

Who this is for. Owners, COOs, CFOs and IT leads of Hong Kong SMEs and mid-sized companies who want a practical first step with AI — one that shows a measurable result within weeks and builds a foundation for the next.

This playbook describes our method. It contains no client names or performance figures; results depend on your data, processes and people.

How to use it

Read chapters 1–2 for the framing. Run the scoring sheet in chapter 3 with the people who own your processes. Use chapter 4 as the plan for the pilot, chapter 5 as the checklist before you start, and chapter 7 to judge whether your organisation is ready. Chapter 6 explains the Azure services you will hear named along the way.

What "AI-native" means in practice

Most business software was designed around forms, records and reports. AI features are then attached: a summarise button, a chat box in the corner. The software still expects a person to find the record, read the report and decide what to do next.

An AI-native solution starts from the other end. The question is not "where can we add AI to this screen?" but "what should happen automatically, and where does a person need to decide?" The design follows from the answer:

- **Workflows have AI in the loop from the start.** Routine steps — matching, extracting, drafting, classifying, forecasting — are done by the system; people handle exceptions and decisions.
- **Data comes first.** The solution is only as good as the data it sees, so connecting, cleaning and governing data is part of the build, not an afterthought.
- **It learns.** Models are measured against reality, corrected by users and retrained. The system that went live is not the system that runs a year later; it is better.
- **It is secure and accountable by design.** Identity, access, logging and content controls are built in, because the system now acts, not just records.

In practical terms, this shapes how a project is run. The first week is spent on data sources and measures. The pilot group uses the solution for real work. Success is a number that moved, not a demo that impressed.

Why bolt-on AI features disappoint

Companies that try an AI feature inside an existing tool and conclude "AI does not work for us" have usually hit one of four problems. None of them is about the model.

The feature sees the wrong data

A chat box attached to one system only knows that system. The answer to the customer's question lives in three others. The assistant is confidently incomplete.

Nobody owns the content

The assistant answers from documents nobody maintains. Within months it is quoting a superseded price list, and staff stop trusting it.

There is no measure

"It seems helpful" is not a result. Without a baseline and a number to move, the feature cannot be defended at budget time and cannot be improved.

Workflow unchanged

If the process around the feature is the same, people use it occasionally and it saves little. The gains come when the routine step is removed from the person's day, not made slightly faster.

The pattern: AI-native projects fix these four in the first week — connect the right data, name a content owner, agree the measure, redesign the step — and only then build.

Choosing a first use case: the scoring sheet

List every AI idea your team has. Then, in a workshop with the people who own each process, score every candidate from 1 (low) to 5 (high) against six criteria. Be honest; the point is to find the project that will succeed first, not the most exciting one.

Criterion	What a 5 looks like	What a 1 looks like
Frequency	Done hundreds of times a week by several people	Done occasionally by one person
Pain	Staff complain about it; it causes errors or delays customers notice	Tolerable; nobody minds
Data availability	Data is in a system, in a usable format, and accessible within the pilot	Data is in people's heads or personal chat apps
Measurability	A number can be stated now: minutes per item, response time, accuracy, forecast error	Benefit can only be described as "better"
Low risk if wrong	Output is reviewed by a person before it matters	Output is applied automatically to customers or money
Ownership	A manager wants it and will make the team use it during the pilot	Suggested by IT; no business owner

Candidate use case	Freq.	Pain	Data	Measure	Low risk	Owner	Total

What usually wins

Scored honestly, the winners are often unglamorous: an assistant for the most common customer or staff enquiries, document extraction for a high-volume form, a demand forecast for one product family, a bilingual draft generator for marketing or listings. These are also the projects that build the foundation — data connections, access model, measurement habit — that the ambitious ideas will need later.

Typical first use cases by industry

Retail & luxury	Clienteling assistant grounded in the product catalogue; bilingual campaign drafting; store-level demand forecast
Finance & insurance	Knowledge assistant over approved product documents with cited answers; extraction from application and claims forms
Logistics & trading	Order-status assistant on WhatsApp; extraction from shipping documents; replenishment forecast for one product family
Professional services	Internal precedent and policy assistant with matter-level permissions; engagement WIP and receivables view
Property	Leasing enquiry assistant in three languages; bilingual listing generation; arrears and occupancy dashboard

The four-to-six-week pilot

A pilot is not a prototype. It runs on your own Azure subscription, with real data, used by real staff for real work, measured against the baseline you agreed. At the end there is an evidence-based decision to scale or to stop — and both are acceptable outcomes.

When	Focus	What happens
Week 1	Scope and sources	Confirm the measure and baseline. Gather the data or documents and name an owner for each. Write the acceptance test (for an assistant: ~50 real questions with correct answers and sources). Settle access, privacy and retention.
Weeks 2–3	Build	Connect data inside your Azure tenant. Build the assistant, model or automation. Run it against the acceptance test; iterate with the business owner weekly.
Weeks 4–5	Real use	A named group uses it in daily work alongside the current process. Track the measure. Collect corrections; fix what is wrong; add what is missing.
Week 6	Decide	Present the measured result against baseline, the cost to run it, risks observed, and a plan to scale (more users, channels, data) or to stop.

Roles

- **Business owner (you):** chooses the use case, provides the pilot group, signs off the acceptance test and the decision.
- **Content or data owner (you):** confirms sources are current and correct; answers questions about them.
- **IT or security contact (you):** grants access in Azure and Microsoft 365; reviews the access model.
- **Delivery team (us):** architecture, build, testing, measurement, documentation and handover.

What you have at the end

A working solution in your tenant, the measured result, documentation of data sources and controls, and a costed plan. If you scale, the pilot becomes the first increment rather than being thrown away.

Data, privacy and PDPO questions to settle first

Hong Kong's Personal Data (Privacy) Ordinance applies to personal data whether it is processed by a person or a model, and the Privacy Commissioner has published guidance on the use of AI. None of this prevents a project; it is a checklist for week one. Settle these with your business owner, IT contact and, where needed, legal adviser.

- ☐ **Which data does the solution need — and which personal data can be left out?** Often the task works without names or identifiers in the prompt at all.
- ☐ **Where will it be processed and stored?** Azure lets you choose the region; your tenant's identity, network and logging controls apply.
- ☐ **Is the use within the purpose the data was collected for?** If not, what notice or consent is needed?
- ☐ **Who can access outputs?** An internal assistant must respect the same permissions as the underlying documents.
- ☐ **How long are prompts, outputs and logs retained, and who can read them?**
- ☐ **Is a person reviewing outputs that affect customers or staff?** Keep a human in the loop for the first project.
- ☐ **What do users and customers get told?** A short notice that an AI assistant is in use, and how to reach a person, is usually enough.
- ☐ **How are errors reported and corrected?** A simple feedback path inside the tool, reviewed weekly during the pilot.

On Azure OpenAI and company data

Azure OpenAI runs generative models as a service inside Microsoft Azure under Microsoft's enterprise terms: prompts and outputs are not used to train the underlying models and are not shared with other customers or with OpenAI; access is governed by your own Azure identity and network controls. Microsoft publishes the precise data-handling behaviour, including abuse-monitoring retention; review it so the statement in your AI policy is accurate. A one-page internal AI policy — which tools are approved for which data, human review before reliance, who to ask — is a sensible companion to the first project.

The Azure building blocks

You do not need to know these services to run a project, but you will hear them named. This is what each one does in plain language, and where it appears in our solution lines.

Service	What it does	Where we use it
Azure OpenAI	Generative models (text, chat, embeddings) as a managed Azure service under enterprise terms	AI Assistant, content generation, natural-language reporting
Azure AI Search	Indexes your documents so an assistant can retrieve the relevant passages and cite them (RAG)	AI Assistant, knowledge bases
Azure AI Document Intelligence	Extracts text, tables and fields from PDFs, scans and photos	Finance & CRM, logistics documents, claims and applications
Azure AI Speech and Language	Transcription, translation, sentiment and topic detection in Cantonese, English and Mandarin	Real-time analytics, assistants on voice channels
Microsoft Fabric and Power BI	Unified data platform and dashboards	Finance, ERP forecasting, analytics
Azure Machine Learning	Training, deploying and monitoring forecasting and scoring models	AI-ERP, churn and lead scoring
Microsoft Teams, Bot Service, messaging	Channels through which staff and customers reach an assistant, including WhatsApp	AI Assistant
Microsoft 365 Copilot	Per-seat AI inside Word, Excel, Outlook and Teams, inheriting your tenant's permissions	Individual productivity; provisioned through our cloud subscription service
Azure DevOps, Functions, API Management	Pipelines, serverless code and secure APIs that connect new solutions to existing systems	Software Development Platform, integrations

All of the above run inside your own Azure subscription. If you do not yet have one, or want Azure and Microsoft 365 billed locally in Hong Kong dollars, we supply and manage subscriptions through the Microsoft Cloud Solution Provider programme.

Readiness checklist

Tick what is true today. Five or more ticks and you are ready to run a pilot; fewer, and the first step is usually a short data or process clean-up that we can scope with you.

Business

- ☐ We have a list of candidate use cases
- ☐ A manager owns the process and wants this
- ☐ We can name the number that should move
- ☐ A pilot group of 5–15 users is available
- ☐ A decision-maker will attend the week-6 review

Data and systems

- ☐ The needed data is in a system we can access
- ☐ Key documents have a known owner and location
- ☐ We use Microsoft 365 (or can provision a tenant)
- ☐ We have, or can set up, an Azure subscription
- ☐ IT can grant access within the first week

Governance

- ☐ We know which personal data, if any, is involved
- ☐ We have (or will write) a one-page AI usage policy
- ☐ Someone can sign off access and retention decisions
- ☐ We accept that outputs are reviewed by people during the pilot

About ICO BROTHERS LIMITED

ICO BROTHERS LIMITED is a Hong Kong-based, AI-native solutions company and a Microsoft Solutions Partner for Data & AI (Azure). We design and deliver intelligent business solutions on Microsoft Azure AI for retail and luxury, finance and insurance, logistics and trading, professional services and property businesses.

Solution lines

- AI Finance & CRM
- AI Digital Marketing
- Software Development Platform
- AI Assistant
- Sales & Customer Service Real-Time Analytics
- AI-ERP

We also supply and manage Azure, Microsoft 365 and Copilot subscriptions for Hong Kong businesses, with HKD billing, migration and support.

Next step

If you would like to run the scoring exercise with us, or discuss a pilot, use the contact form on our website. We work in Cantonese, English and Mandarin.

ICO BROTHERS LIMITED

Room 602, 6/F, Kai Yue Commercial Building, No. 2C Argyle Street, Mongkok, Kowloon, Hong Kong
Monday – Friday, 9:00 – 18:00 (HKT)