STRATEGY BLUEPRINT

When the team reaches for ChatGPT

Official systems feel slow or empty, so people paste work into ChatGPT or similar tools. Someone bought Copilot or “AI features” on top of existing software.

When the team reaches for ChatGPT

The situation

Official systems feel slow or empty, so people paste work into ChatGPT or similar tools. Someone bought Copilot or “AI features” on top of existing software. Leadership hears about productivity wins in anecdotes and about risk in the same week. There is no shared place where AI assistance is good, governed, and tied to how the business actually runs.

Problems it causes

- Client and internal data leave the boundary into consumer tools you do not control.
- Answers are fluent and sometimes wrong, with no audit trail tied to a job or client record.
- Two people get two different “AI” answers for the same process step.
- Shadow AI creates uneven quality and hidden dependence on individuals’ prompting skill.
- Bolted-on copilots sit beside a process that still does not fit, so handoffs do not improve.
- Security, legal, and insurance conversations get harder because usage is informal.
- You cannot improve the system: prompts and context are not part of an owned workflow.

Why the usual fixes fail

A ban without a better tool drives usage underground. Generic enterprise chat still lacks your stages, roles, and approvals. Vendor AI toggles inherit whatever is wrong with the underlying CRM or ticket tool. An “AI transformation” programme sells workshops while the team keeps pasting into ChatGPT.

The path forward

Keep what still earns its keep. Integrate on purpose. Replace the wrong centre. Useful with or without FluentForward.

01 Admit where shadow AI already happens (which roles, which tasks, which data).

02 Rank tasks: assist (draft, summarise, classify) versus decide (approvals that must stay human).

03 Draw the data boundary: what may never leave your tenancy; what can be processed in a controlled service.

04 Keep consumer ChatGPT out of the core workflow; allow only sanctioned tools for low-risk personal productivity if you must.

05 Keep best-of-breed model APIs as infrastructure if useful, but behind your application, not as the workplace.

06 Replace “paste into chat” with an owned system shaped around the real process (the same spine logic as the other blueprints).

07 Put AI inside specific steps of that system (for example draft a reply from the client record, classify an inbound request, suggest next stage) with the record as context.

08 Log prompts and outputs against the job or client so review and improvement are possible.

09 Integrate identity and permissions so assistance respects the same access rules as the rest of the system.

10 Train the team on the workflow in the system, not on clever prompting in a public chat window.

11 Measure leak risk down (shadow tools unused for core work) and cycle time or quality up on the assisted steps.

12 Grow assisted steps only where the process is clear; do not sprinkle AI on a broken handoff and call it transformation.

What good looks like

Core work stays inside a system you own. AI speeds the steps that deserve it. Data and judgement do not leak into a consumer chat. Quality is consistent because context is the business process, not a one-off paste. Security and legal can describe where assistance runs and what never leaves the tenancy. Quality becomes consistent because context is the job or client record, not a one-off paste. The team trains on the workflow in the system, not on clever prompting in a public chat window.

Example first release (about 8 weeks)

FluentForward's focused delivery shape for this problem. Business milestones, not the full path list.

Weeks 1-2	Map where shadow AI already happens; rank assist versus decide; draw the data boundary for what may never leave your tenancy.
Weeks 3-4	Design assisted steps inside an owned process spine; choose sanctioned model infrastructure behind your app; agree which consumer tools are out for core work.
Weeks 5-7	Build those assisted steps with the record as context; log prompts and outputs against the job; wire identity and permissions.
Week 8	Go-live on the first assisted steps; stop shadow paste for that core work; short adoption note; grow only where the process is clear.

One line a peer can forward

"We stopped pasting client work into ChatGPT and bolted-on copilots, and put the right AI inside a system we own that follows how we already work."