

Developing an AI enabled finance team

Article 4 of 6 · aicfopartner.ai · PDF issued 9 September 2026

EXECUTIVE SUMMARY

- Deploy an AI tool that works straight out of the box, such as Claude, ChatGPT or Microsoft Copilot. Using it on familiar tasks can save time early while building the team's AI fluency.
- Choose a business or enterprise plan and keep task instructions, source files and checks reusable across tools. This makes it easier to change provider without redesigning the underlying work.
- Start the automation journey with a few tasks that are simple to hand over, stable month to month and necessary over the longer term.
- Run a practice phase of about six weeks inside the minimum governance framework. Measure hours, errors and spend against a pre AI baseline for each automated task.
- Save the instructions and checks for proven tasks as reusable Skills, so one person's technique becomes a capability the team owns.
- Bring the team along with empowerment and clear communication: an amnesty for existing use, visible leadership and celebrated wins.

01 — CREATE CAPACITY WHILE BUILDING FLUENCY

Most finance teams already know what they want from AI, more time and better, quicker data. However, teams are often constrained by bandwidth and find successfully adopting AI challenging. As discussed in the strategy article, the answer is therefore not an AI deployment programme that starts by materially adding work, rather an approach designed to remove it quickly. Deferring non priority tasks helps, but the most powerful step forward is deploying an AI tool that is effective straight out of the box. Such tools are extremely intuitive so when a team begins using them on routine tasks, such as variance commentary or personal administrative workflows, task completion times decrease rapidly. Then, because people learn tool functionality fastest on work they already know, the same work that frees the time also builds AI fluency. In summary, the sooner such an AI tool is adopted, the faster the dual effects of time savings and AI proficiency compound.

However, adopting the technology within the team is usually the easy step, the cultural shift needed to get the best out of the investment being the harder part. A busy team has to be persuaded to work differently, people need confidence and support to use unfamiliar tools on routine workflows. Adoption is often done best when the shift feels like an investment in the team's capabilities rather than a threat. A team that feels equipped and trusted will develop quickly and find uses for AI that no plan predicted, while a team that experiences change without being empowered or supported will quietly return to the old way once attention moves elsewhere. Mounting evidence supports this point. KPMG's May 2026 survey of a thousand finance leaders found that active AI use in finance has

more than doubled in two years. Yet the most cited barrier among those unhappy with the returns was change management, and the biggest training obstacles were the lack of role specific use cases and of hands on practice. The good news for a stretched team is that neither of these barriers requires material spend, only a sound approach. This article goes through that approach, outlining the steps that underpin it: tool setup, initial tool deployment, how to manage deployment, how to bring the whole team along and what comes next. One caveat, hopefully an intuitive one, runs through all of it: anything installed should be reviewed with IT before it is enabled.

02 — CHOOSING AND SETTING UP THE AI TOOL

Deciding which AI tool to adopt and setting it up correctly is clearly the foundation from which everything flows. This article covers setup with the Claude ecosystem in mind but happily ChatGPT with Codex and Microsoft Copilot broadly cover the same ground. The principles and basic steps here carry across the different vendor choices, which is more than a convenience. As the strategy article sets out, tool neutrality belongs at the front of the design when integrating AI into the finance function. [No single AI tool stays ahead for long](#). Pricing shifts regularly, capability leapfrogs between vendors, and geographic access can change with little warning, as recent months have shown. A setup that lets the team switch AI tool without rebuilding its work is therefore worth more than any single subscription, and the habits in this guide, the designated folder, the task written down, the checks that make output trustworthy, carry across ecosystems even where the feature names differ.

Whatever AI tool choice is made, a business or enterprise subscription is highly recommended and in many cases will be a non negotiable for IT. At these subscription tiers any information input is not used for model training by default, administrators control data access and retention, and audit logs and spend caps are built in at Enterprise tier level. Once a subscription choice is made, users are decided and assigned, and the applications can be installed on company devices. For Claude users this means the app, which includes chat and Cowork, plus the Excel and PowerPoint add ins. Cowork should then be granted access to relevant and appropriate finance folders and nothing beyond.

Once at this stage, relevant connectors, which are integrations with existing commercial software applications, can be reviewed and linked to Claude. These are accessed by clicking through to connectors in settings within the Claude application. The list can simply be browsed for what is applicable to the organisation. Outlook, OneDrive, SharePoint and Gmail are a good place to start. There are also connectors into a growing number of ERPs and finance adjacent applications, NetSuite and Xero among them. Each connector authorises per user, so Claude sees only what that person can already see.

The next area to look at is plugins, which are pre built sets of Skills that Claude can read and learn from, and which can save teams time, especially initially. Anthropic, for example, publishes [an open source finance plugin for Cowork](#), covering journal entry preparation, reconciliations, variance analysis and close management. [Its financial services marketplace, launched in May 2026](#), adds agent templates split between analyst and operations work. The ledger reconciler and month end closer translate well to a corporate close, once team AI skills have developed to become familiar with agentic workflows. Useful as these plugins are though, the lasting value tends to come from bespoke Skills the team develops, which essentially become a process record for the finance function, as discussed in more detail in the next section.

At this stage, it is useful to mention the companion resource to this section, [the Quick deploy AI enabled finance tasks page in the Resources section](#). This resource carries the practical side of this introductory approach in one place, covering the following areas:

- Which Claude surface (chat, Cowork, add ins) suits which kind of task.
- How Cowork works with files and folders.
- Data safeguards and initial governance setup.
- Prompt frameworks that make every output checkable.
- Twenty nine worked finance task examples across payables, receivables, close, FP&A, contracts and admin, each with the steps and the checks that make the output trustworthy.

This article covers the decisions and the sequencing behind initial AI deployment, and the Quick Deploy resource is the working companion the team can refer to when they sit down to build a workflow in Claude.

03 — CHOOSING THE FIRST TASKS

The discipline that matters most at the start is focus. Many finance teams starting with AI face the same temptation: a long list of tasks to automate and a desire to approach them all immediately. Teams that spread themselves across the list tend to experience inconsistent execution and learning outcomes. The best approach is to start with three or four suitable tasks based on the following criteria:

- Can the task be explained to a colleague with a short verbal handover? Work that can be explained in five minutes can be described to an AI tool with the same ease.
- Is the task stable, drawing on the same sources, fields and exceptions month after month? A task that changes shape with each cycle means the time saved in one place is lost in another.
- Does the task deserve to exist at all? The selection process is a natural moment to notice work that nobody would miss, and that work is better left unautomated than automated.

It is also recommended that teams begin with FP&A tasks, where learning is fastest and the work can be kept read only. This means mistakes do not affect the underlying transaction data and are typically caught in standard review steps.

Transactional finance then follows as confidence and AI skills grow, and the opportunity here is usually the biggest. The work is high volume and repetitive, so the recoverable hours add up quickly.

This initial focused approach, built around a small number of tasks, will typically produce quick results in an engaged finance team. Once an AI based task is running reliably, it is worth capturing exactly how it is run before the method starts to vary from person to person. In the Claude ecosystem, the mechanism for doing this is a Skill. A Skill in Claude is simply the documented task steps, that is the prompt that runs it, the steps it follows, the files it draws on, and the checks it has to pass. Each skill is then stored once in the shared workspace. From then on, anyone in the team can run the task with a one line request rather than working it out afresh. What began as one person's private technique becomes a capability the whole team owns. Each Skill should have a named owner, so that any later change to it is a decision somebody takes rather than a drift nobody notices. A common way to build a Skill is to simply ask Claude to do so once it has correctly completed the task. The task owner then reviews and validates, saving the Skill once correct.

04 — RUNNING THE PRACTICE PHASE

A supportive context matters just as much as task selection, especially in the first months of AI adoption. The recommended approach here is a practice and innovation phase, deliberately initiated and clearly communicated. The team should know what the coming weeks are for and that early stumbles carry no penalty. Indeed, a strong feedback loop to learn from successes and failures is critical to success. This initial phase works best with a firm end date, around six weeks out, scheduled away from busy periods such as the audit, month end and budget season. An initial phase without a deadline tends to drift into open ended experimentation, always interesting and never concluded, with no tangible benefit. Initial governance for the phase stays deliberately light and fits [the criteria outlined in the governance article](#). The key steps are an approved AI tools list, a person in the loop on anything touching the general ledger, a record of which tasks may run live and review protocols, and a named owner for every automated task.

How the team learns also needs consideration. There is a place, of course, for structured learning, especially overviews of tool functionality. However, most people tend to learn how to use new technologies best by watching a credible colleague do real work in front of them. Training therefore works better as a demonstration than as a presentation. One person shows a working task using real data, and the others adopt the approach where it applies to their own work. Whoever is furthest ahead runs the sessions, whatever their grade. Then of course there is no substitute for practising what has been demonstrated.

Another facet of successful deployment is keeping measurement and ROI firmly in mind throughout the practice phase, and on an ongoing basis. Before the AI tools are deployed, record how long each task takes today, from first touch to finished output. Every later claim of time saved is then judged against a fixed starting point rather than a shifting memory. Count total process time rather than drafting time. The error rate also needs tracking alongside the hours, assessed by simple measures such as counting the corrections a reviewer makes to each draft. After all, an hour saved in drafting that reappears as an hour of correction has not been saved, only moved along the process. When genuine hour savings are realised, they should be reassigned in writing to specific tasks such as business partnering or process improvement. Freed time that is not formally reallocated is absorbed back into the day to day workload, and the gain quietly disappears.

Cost is the other half of the same ledger. Seat plus usage pricing is now common for these tools, so the spend meter runs from the first day of the practice phase, and ungoverned AI spend is a finance discipline failure that lands on the CFO. The guardrails belong at adoption rather than after the first surprising invoice: spend caps set from day one, a weekly usage review during rollout, and each person's use monitored as management information, serving as cost control and adoption signal at once. Together the two sides give the phase an accurate return: hours saved and error rates held on one side, the metered cost of achieving them on the other.

Finally, as alluded to in the introduction, capacity, or the lack of it, is often the biggest barrier to success. Deploying an AI tool within a well run learning period creates a virtuous cycle of efficiency and AI skill development. However, it is not always enough. Some finance teams, especially those running complex processes across a patchwork of Excel and assorted applications, need a further boost.

A strong early move in this respect is a watch and learn tool such as GoApprentice, which automates tasks that can be demonstrated on a user's desktop while giving the team a feel for how an automated workflow operates. The vendor says it learns a finance task by observing it once, then sets up an automated workflow to replay it at the requisite cadence. The vendor states SOC 2 and ISO 27001 certification, the sensible minimum for any tool

watching finance work. It also has the added advantage of generating a standard operating procedure during the learning process, verified and iterated with the process owner until complete. Further detail sits in [the GoApprentice entry in the Resources section](#).

05 — BRINGING THE TEAM ALONG

Beyond the technical and tactical, successful AI deployment is fundamentally about bringing the finance team along, through empowerment, enjoyment and clear communication. The strategy article describes the middle layer barrier, which manifests when controllers and team leads are expected to carry the weight of AI integration and process redesign [on top of the day job](#). The approaches outlined in this article help solve this issue from a tactical perspective, but this on its own is not enough. Two further things keep the middle moving:

- Adoption needs a working owner, usually the CFO, who uses the tools personally, picks the first tasks and gives cover while results lag effort.
- The managers in the middle then need a short written brief that sets out what is in scope this quarter, what stops or moves to make room, what may change without approval, and where to go when something is stuck.

It is also worth bearing in mind at this stage that most teams will arrive with their own AI habits. PagerDuty's June 2026 survey of 1,250 office professionals found that two thirds had used AI at work, despite believing it was not permitted, and there is little reason to assume a finance team is different. The answer is an amnesty and a simple use log rather than a ban. A ban pushes use underground, where the risk is invisible, while the log can be used to consolidate tools and spend into an approved list.

The change should also be enjoyable, and that is not a soft point, because adoption spreads through a team on energy as much as on instruction. A shared channel where colleagues post the problem, the prompt, and the results costs nothing to run. A twenty minute demonstration slot on rotation weekly can keep the examples fresh. Routinely scheduled hackathons after the practice phase surface enthusiasm, innovation and talent together. A leaderboard can also help, if it ranks wins, hours returned, and Skills shipped, not simply token consumption, which rewards spend without measuring value. Recognition tends to beat prizes, though the latter can certainly help: wins celebrated publicly, AI use written into objectives and appraisals, and the leader seen using the tools in their own board pack are all worth putting in place.

06 — MOVING BEYOND THE PRACTICE PHASE

This initial practice phase will almost always surface an uncomfortable discovery or two, because working closely with a process is the fastest way to notice what is wrong with it. The natural instinct at that point is to pause and redesign before automating anything further but that instinct can be mistaken. Provided a process is in reasonable shape, meaning the steps are known, the inputs are reliable and the output can be checked, automating it improves matters straight away, even where the process itself could clearly be better. The hours come back immediately, and that new bandwidth is what lets the team optimise the process in question or turn attention to areas of opportunity elsewhere in the function. As the strategy article argues, more programmes die waiting for the flawless dataset or the fully redesigned process [than from any technology choice](#). The philosophy at this stage should be that good enough beats perfect.

Beyond that point the questions change shape. Some tasks outgrow the chat and Cowork patterns and evolve towards proper workflows, chains of steps across systems that deserve orchestration rather than a saved prompt or Skill. Some solutions are better bought than built: the existing ERP may already carry more AI capability than the current contract tier exposes, and a market of specialist tools now covers the high volume ground in payables, receivables and reconciliation. Conversely, some capability is worth assembling precisely because it is deployed on a cost effective and bespoke basis. Weighing those routes properly to understand what to buy, assemble or build and in which order, is [where this series goes next](#).

THE AI IN FINANCE SERIES

- 01 [AI in finance: what it is and how it works](#) [aicfopartner.ai/ai-in-finance](#)
- 02 [AI strategy for finance: the six choices](#) [aicfopartner.ai/ai-strategy-for-finance](#)
- 03 [The roadmap to the AI enabled finance function](#) [aicfopartner.ai/ai-finance-roadmap](#)
- 04 [Developing an AI enabled finance team](#) [aicfopartner.ai/ai-finance-team-development](#)
- 05 [Sourcing AI capability: buy, build or assemble](#) [aicfopartner.ai/ai-buy-build-or-assemble](#)
- 06 [Governance that speeds AI adoption](#) [aicfopartner.ai/ai-governance-for-finance](#)

Questions or comments: tim@aicfopartner.ai