AI integration services: what to agree before AI touches your systems

By Andy Slack | November 7, 2026 | AI Development

AI integration services, explained for buyers: what the AI reads, what it writes, where a person approves, and what happens when it can't.

An AI integration puts a model inside software and work your business already runs. It reads something from one system and leaves a result in another, where your team already works. Before a partner builds one, agree five things: what the AI reads, what it writes, where a person approves, what happens when it can't settle something, and what your team keeps at the end. This article is those five, with one of our own integrations read against them.

Two neighbouring questions have their own pages. Which process to start with is in what to automate first, and when to build (/articles/business-process-automation-what-to-automate-first). Whether the job needs a connector, shared tasks or software written for it is in workflow automation software (/articles/workflow-automation-software).

What an AI integration is, and what it isn't

An integration works with the systems you already have. When the system itself has to be built, that is bespoke development (/bespoke-development): custom software built around the workflows that matter to you. When the work is your engineering team's own delivery, such as coding agents and AI development tooling, that is AI-assisted development (/ai-assisted-development). This page is about the first case: AI added to an operation that already runs.

Five things to agree before anything is built

- What the AI reads. One kind of input from one place: an attached document, a record, a message. A narrow input is easier to test against your own examples.
- What the AI writes, and where. Agree whether it fills in a draft or takes the final action. A draft that a person checks is the safer first output.
- Where a person approves. Put the approval in a tool your team already uses, so nobody has to go looking for it.
- What happens when the AI can't settle it. Every integration needs an exception path: who gets the case, and with what attached.
- What your team keeps. The training and the written record that let your team run and change the integration without the partner.

If a proposal can't answer all five in plain words, it isn't ready to build.

One of our integrations, read against the list

For one client, we automated finance admin (our project record, confirmed 7 November 2026). Read against the five:

- Reads: the invoice PDF attached to a bill.
- Writes: the missing fields on the draft bill.
- Approves: bills that need approval go to Slack, with a one-click approve button.
- Can't settle: anything the AI can't settle goes to a person, with the document attached.

The AI step fills in a draft, and the approval is a button in Slack. That boundary is the part worth copying: the model's job is narrow enough to check, and the approval sits outside the AI step.

We first mentioned this work in what to automate first (/articles/business-process-automation-what-to-automate-first), where the point was picking the process. Here the point is the boundary around the AI step. Two things are left out on purpose. The accounting system and the connector aren't named, because the record we can publish doesn't name them. And there is no outcome figure: the results we publish are for a sector, and they don't show what this one workflow did.

Questions to ask about access, errors, tests and monitoring

Put these to any partner, us included, and ask for the answers in writing for your own integration:

- Access. Which systems will the integration read, which can it write to, and under whose account?
- Errors. When the AI's answer is incomplete, or the system it writes to is down, what happens, and who finds out?
- Tests. How will it be checked against your own documents before it goes live, and again after each change?
- Monitoring. How will you know it has stopped working before a supplier or a customer tells you?

For the finance work above, the exception path is the part our record shows: anything unsettled goes to a person with the document. We don't claim particular access controls, tests or monitoring for it here, because they aren't in the record we can publish.

Where the AI step sits

The AI step can live in a workflow tool, in shared tasks with an approval gate, or in software written for the job. We set out that choice in workflow automation software (/articles/workflow-automation-software) and don't repeat it here. If the integration grows into a product your own customers log into, read a SaaS development company ships the product, not a directory (/articles/saas-development-company).

What your team keeps

Our general handover practice is set out in what to automate first (/articles/business-process-automation-what-to-automate-first): full training, so your team can run and change what we built, and a written summary of what was built and how it runs.

For an AI integration, ask whoever you hire for that summary to cover the five things above. Then your team can change what the AI reads, what it writes and who approves without calling the partner back. That is what no vendor lock-in means in practice.

If you want a second view before you choose a partner, our consulting (/consulting) covers AI and automation strategy. If AI is going into how your engineering team delivers, start with AI adoption (/consulting/ai-adoption). There is more in our articles (/articles).