

What is a human agentic workflow?

A plain-language guide to AI that does the steps while people make the calls

Audience	Business owners and managers
Version	0.1
Published	October 21, 2026
Online	omnafy.com/whitepapers/human-agentic-workflows/

Abstract

Most office work is a mix of two things: steps that are the same every time, and decisions that need a person. A human agentic workflow splits them on purpose. AI handles the repeatable steps, and a person approves the ones that carry consequences. This paper explains what that means in practice, walks through one example from start to finish, and is plain about what this approach is not good at yet.

1. The short version

Picture the work that lands on one person in your office every day. A form comes in. They read it, type the details into one system, then into another, notice something is missing, and send an email asking for it. Somewhere in the middle there's a real decision: approve this, flag that, call the customer.

Almost all of that time goes to the steps, not the decision. A human agentic workflow hands the steps to AI and keeps the decision with the person. The person still sees everything, still signs off, and can still say no. They just stop retyping.

2. What an agent actually does

"Agent" is the word the industry uses for AI software that can take actions instead of only answering questions. In plain terms, an agent can:

- Read documents, emails, and forms, including scans and photos.
- Look things up in the software you already use.
- Fill in records and draft messages.
- Follow a set of steps you've agreed on, in order.

What makes this different from older automation is the reading. Traditional automation breaks the moment a form arrives in a slightly different layout. An agent can read a messy fax or a photo of a field ticket and pull out what matters, the way a person would.

What an agent shouldn't do is decide things that a person is accountable for. That's the line this whole approach is built around.

3. Where people come in

Every step in a workflow gets a level that says how much a person is involved. Looking something up runs on its own. Changing a record inside your systems might be checked before or after. Anything a customer will see waits for approval. Judgment calls always belong to a person.

The approval itself has to be quick to do well. A person reviewing an agent's work should see what it did, what it based that on, and anything it wasn't sure about, all in one place. If they have to open three other systems to check, they'll stop checking. Our companion paper, *Where a person should sign off*, covers this in detail.

4. A worked example

Here is how it looks for an insurance agency handling certificate of insurance requests.

Something comes in. A contractor's client emails asking for a certificate naming them as certificate holder.

The agent does the legwork. It reads the email, finds the contractor's policy in the agency management system, and drafts the certificate. It notices that the request asks for additional insured wording the policy doesn't include, and flags it.

A person signs off. The CSR opens a short summary showing the draft certificate, the policy it came from, and the flagged wording. They approve the certificate and take the flagged question to the underwriter.

The CSR made every decision that mattered. What they skipped was the lookup and the typing.

5. What it isn't good at

We'd rather say this up front:

- **Work with no pattern.** If every case is truly different, there's little for an agent to repeat. These workflows pay off on work that happens often and looks similar each time.
- **Processes nobody has written down.** If the steps live in one person's head, the first job is getting them out. That's useful on its own, but it takes time.
- **Decisions that need judgment.** An agent can gather facts and lay out options. It shouldn't make the call, and we don't design workflows where it does.

Omnafy is early. We design each workflow case by case, and we haven't published measured results yet. When we do, they'll come from real workflows, not estimates.

6. How we build one

We start by sitting with the people who do the work today and picking one task worth handing off. We map every step and agree on which ones an agent handles and which ones need a person. You approve that plan before we build anything.

Then we connect the workflow to the software you already use and run it alongside your team. Every step starts with a person checking it. As the agent proves itself on a step, you can choose to let that step run on its own, and you can put the check back at any time.