



MAKES ELEVEN CONSULTING · THE METHOD WE RUN

Audit. Gauge. Engineer. Navigate. Track.

Five moves that take a workflow from “this is slow and nobody knows why” to “this is measurably better and the CEO believes the number.”

THE A.G.E.N.T. METHOD, A SHORT READ

Every engagement runs the five in this order, and each stage ends in an artifact you keep. Four questions come before the first letter. One engine runs under all five stages, and the published sprint is roughly eight weeks. This short read sets out all of it, with a worked example on a real workflow.

The five stages

Audit · Gauge · Engineer · Navigate ·
Track

Before them

Four questions, and an afternoon

Under them

Signal · Enrich · Score · Compose ·
Act

A Guardian ships on every engagement, the examiner is never the builder, and no agent signs, files, or pays. The full enumeration of what our agents refuse is a companion document, the prohibited-actions register.

THE METHOD, STAGE BY STAGE

Four questions first, then five moves in order.

Before the first letter

Four questions come before Audit. They begin with what the business is for rather than with a workflow, and answering all four takes an afternoon, not a quarter.

Business objectives. What the firm is trying to do this year; everything downstream traces back to it.

Your agentic position. Where machines act, where they must not, and where the refusals get written.

Opportunity identification. Across the value chain, systematically, never from a vendor list.

Select the first workflow. Strategically important, but safe enough to experiment on. One, chosen deliberately.


The five stages

A Audit	<i>"What is actually happening, and why does it fail?"</i> Understand how the work is done today and what outcome really matters, rather than what the process document says. You keep. A written picture of the as-is and the target outcome. Useful even if you stop here.
G Gauge	<i>"What would success look like for the business?"</i> Score each workflow on impact, repeatability, complexity and outcome potential. This is where we decide what not to build. You keep. A prioritized shortlist of agent opportunities, and a documented list of what we declined to touch.
E Engineer	<i>"If you started fresh with AI, what would this look like?"</i> Redesign and build the flow so an agent can act: data accessible, decisions explicit, guardrails shipped with it. You keep. A working agent-first process ready to pilot, with measurable success criteria attached.
N Navigate	<i>"Who decides what, and what happens when AI is wrong?"</i> Shape the human-agent relationship: transparency, intervention paths, governance. The system prepares expert work for approval. You keep. Human control that holds as autonomy rises, because the escalation path was built rather than promised.
T Track	<i>"Would your CEO believe these numbers?"</i> Prove value against the outcome chosen in Gauge, measured against what would have happened anyway rather than simply before and after. You keep. Outcome gains that fund the next move, or a documented reason to stop. Both are results.

UNDER EVERY STAGE

One engine. Five moves. Two things change.

Every agent we ship runs the same five moves. Only two things ever change between a legal docket watcher and a bookkeeping close agent: what stage one watches, and what stage five may do. Every engagement hardens the same spine.




A worked example, on a real workflow. Signal: a recertification falls due in 45 days. Enrich: the file assembles itself from the tenant file, the prior certification and the agency rule. Score: every claim gets a weight, and a claim that falls below the threshold becomes a question to a human rather than a sentence. Compose: a draft that shows its sources, line by line. Act, only as far as the gate allows: the filing is sent by a person, never by the agent.

THE SHAPE OF THE WORK

One workflow. Roughly eight weeks. Then you decide again.

The published sprint runs in three phases. We keep that shape because it forces a visible result early, while the people who will live with the system are still in the room.



Scope and duration are set per engagement: a six-person practice and a national rollout are never the same shape.

The A.G.E.N.T. method, its agent taxonomy and its sprint structure are adapted with attribution from Hofmann & Kruhse-Lehtonen, "The agent-centric enterprise," *Harvard Data Science Review*. What the firm adds is the assurance architecture around it: a Guardian on every engagement, the examiner never the builder, sourced answers or a refusal, and a written register of the actions our agents refuse. No agent signs, files, or pays.