




Helmsted

THE RIA'S CHARTBOOK · NO. 03


FOR INDEPENDENT RIAs · PROACTIVE ADVICE

Faster Check-ins, or Better Health?

Why the real promise of AI for independent RIAs is proactive care for the household, and the harness that makes it trustworthy.

INSIDE THIS GUIDE

- 01 Efficiency is not the destination
- 02 Why AI needs a harness
- 03 Trust is earned the way it always has been
- 04 Define what AI may do with the information
- 05 Walk one client decision through to completion



SECTION 01

Efficiency is not the destination

Picture two medical practices adopting new technology. The first uses it to shorten the waiting room: faster check-in, automated reminders, a smoother front desk. Patients notice, and the practice runs a little leaner.

The second uses it to keep watching the patient between visits. It follows the numbers that matter, notices when something drifts, and brings the physician the right question before a small issue becomes a serious one. The visit itself changes, because the physician arrives already knowing what needs attention.

Both practices “use technology.” The second makes ongoing outcomes its organizing principle.

Independent RIAs face the same choice with AI. It is tempting to sprinkle AI on the legacy workflow: summarize the meeting faster, draft the email sooner, fill the form with fewer clicks. Those gains are real, but they are the waiting-room version of the opportunity. The larger one is proactive care for the household: continuous attention to a family’s accounts, plans, deadlines, and changing circumstances, so the adviser can act before a problem arrives and spend more time on the decisions that build wealth.

That is also where AI moves from answering questions to doing work. The adviser and client decide what should happen. AI carries the authorized steps through to completion, surfaces exceptions, and brings back evidence of the result.

The adviser and client decide. AI carries the authorized work forward.

There is a condition attached. Proactive care runs on information. A system that watches a household continuously needs broader access to client data than a tool that drafts one email, and a system that acts needs clear authority. The more ambitious the use of AI, the more the protection of client information becomes the foundation rather than a feature.

For an advisory firm, the useful question becomes: **Can this system carry out the work we intend, with the information and authority we intend to give it, and can we prove how that information is protected?**

No. 01 covered protecting the information; No. 02 covered the evidence that earns a system permission to act. This installment describes the destination: what the adviser's work looks like when AI carries the execution, and the harness that makes it trustworthy. The workflow recommendations are editorial guidance.



The adviser and client decide what should happen. AI can carry the authorized work forward.


PART II · THE HARNESS

Delegation that knows its boundaries.

SECTION 02

Why AI needs a harness

A language model produces its work step by step, each step conditioned on what came before. NIST notes that this “can produce factually accurate and consistent outputs,” and also outputs that are inaccurate or inconsistent, often stated with confidence. [1, section 2.2] Think of a ball rolling across a table: over a short distance its path is predictable, but over a long one small drift compounds. Research on multi-step reasoning finds a similar pattern, with performance that “can rapidly decay with increased task complexity.” [2] Reliability in practice depends less on the model alone than on what surrounds it: the information it works from, the tools it may use, and whether each step is checked. That is not a reason to avoid AI. It is a reason for a harness.



Each step follows the last. The harness keeps the path close to where it was meant to go.

A harness is everything that keeps each step close to the intended path:

PART OF THE HARNESS	WHAT IT DOES
Grounding	The model works from verified household records and approved sources, not from what it happens to recall.
Scope and instructions	Each task has a defined purpose, household, and output, so there is less room to wander.
Permissions	Tools and systems are limited to what the task needs; the model cannot reach what it has not been granted.
Checkpoints	Results are checked against their sources and rules before the next step builds on them.
Decision points	Consequential steps return to the adviser; exceptions stop the work rather than guessing.
The record	Every instruction, action, and result is logged, so the path can be reviewed afterward.

Guidance on building AI agents makes the same points: an agent should gain “ground truth” from its environment at each step, pause for human feedback at checkpoints or when it meets a blocker, and run with stopping conditions such as a maximum number of iterations. [5]

Notice that the harness and client-data protection are the same structure. What the model may see, what it may do, and what gets recorded are security questions and quality questions at once. A well-built harness produces better results and a more defensible data trail.

It also changes how to evaluate technology. Models improve quickly and are increasingly interchangeable. The harness is what determines whether AI can be trusted with a household’s information and with real work. When comparing options, ask less about which model sits underneath and more about how the system grounds, limits, checks, and records what the model does.

Trust is earned the way it always has been

Eventually, people will trust AI to act for them. Advisers will delegate more of the firm's work to it. Clients will too: many will arrive with their own AI assistants, asking questions, requesting documents, and expecting answers on their schedule. Professionals should start thinking now about what that means for their practice, rather than assuming today's tools will carry them through.

That is where a harness matters most. It is how a firm learns what AI can reliably do: start with bounded tasks, check the results against their sources, widen authority where the record supports it, and keep the adviser's judgment where it belongs.

Delegation tends to widen in stages:

STAGE	HOW THE WORK RUNS	ADVISER EXPERIENCE
Assist	AI prepares the work; the adviser reviews and approves each step.	Operates screens: opens each system, checks each draft, clicks through each step.
Delegate	AI carries routine work under standing authority; the adviser approves by exception.	Approves proposed work: reviews a queue of prepared actions and the evidence behind them.
Operate	AI runs the household's routine operations end to end; the adviser sets direction, makes the judgment calls, and audits the record.	Directs outcomes and handles exceptions: states the goal, decides what needs judgment, reviews results.

Most firms are in the first stage today. The harness is what lets them move through the others.

Early evidence from one domain shows how oversight changes as trust is earned. A study of a coding agent in software development found that experienced users shifted from approving each action toward monitoring and intervening when it matters: new users ran with full auto-approve in about 20 percent of sessions, experienced users in over 40 percent, and experienced users also interrupted more often when something needed attention. [6] That is software development, not financial advice, and it does not show that graphical interfaces will disappear. It does show the direction: from approving every step to stepping in where judgment is needed.

In Helmsted's vision, advisers express intent, make decisions, and resolve exceptions. AI becomes the execution layer, carrying authorized work through to a verified result. Standard point-and-click workflows give way to interfaces organized around decisions and outcomes. The controls stay visible: the adviser can review the evidence behind any result, change the authority a workflow holds, and stop work at any time.

Trust in technology has always been built this way. Few professionals re-add a spreadsheet column by hand anymore. The calculation is not trusted because someone insisted on it; it is trusted because it is consistent, testable, and easy to audit. Pilots rely on autopilot every day, within defined modes, with the crew monitoring and clear rules for when to take over. Nobody trusted either on day one. Structure, verification, and a track record earned it.

AI will follow the same path. The firms that build the harness now, and learn with it, will be the ones clients trust when delegation becomes the norm.

SECTION 04

Define what AI may do with the information

As systems gain the ability to execute work, access and authority need to be considered together.

An application may need permission to read a household record without permission to change it. A workflow may be allowed to prepare a message without being allowed to send it. An approved decision may permit a defined sequence of administrative actions while still requiring a fresh decision if the facts change.

We recommend describing that authority at the level of the task. AI might **prepare** a meeting brief from approved sources for a specified household; **maintain** designated internal fields when a validated event occurs, keeping a history; **coordinate** tasks and approved routine reminders to verified recipients; **carry out a decision** through an agreed administrative sequence within named systems, recipients, and limits; and **escalate** when instructions conflict, a material fact changes, a recipient is uncertain, or an action exceeds the authorization.

This is a proposed operating model, not a description of any product's currently available features or a grant of trading, money-movement, or legal authority. Transactions and other consequential actions require the permissions and controls applicable to those activities.

The objective is to make delegation useful: standing authority for a suitable recurring task, with a clear scope, an owner, a way to withdraw it, and a process for exceptions.

Before granting it, the firm should test whether the workflow stays inside its boundary, using the execution evidence described in No. 02: blocked unauthorized actions, verified completion, and a clear way to revoke authority. “Completed” should mean that the intended destination confirmed the outcome, with unresolved steps made visible.

SECTION 05

Walk one client decision through to completion

Return to the Morgans, the fictional couple three years from retirement who appear throughout this series. In No. 01 their records were mapped and protected; in No. 02 the system earned permission to request their documents. In the waiting-room version, AI helps the adviser write the meeting recap faster. In the proactive version, the work began earlier: the system noticed that a projected income gap had widened after a change in the Morgans’ accounts, and brought it to the adviser before they asked.

The adviser and the Morgans meet, work through the implications, and agree that the next step is to update the income analysis and coordinate with their CPA before making an implementation decision.

In a well-defined future workflow, AI could assemble the approved records, refresh the analysis, identify missing information, create the follow-up tasks, and prepare the material for the CPA. A permitted routine request could go to a verified recipient if the firm has established the necessary authority and disclosure basis. An unexpected recipient or a request to transmit additional documents would trigger review.

The adviser evaluates the planning implications. The client makes the relevant choices with the adviser. Once the next steps are authorized, the system can carry the permitted work forward and report what was completed, what failed, and what still requires a decision.

ONE DECISION, TWO LANES

THE MORGANS · ILLUSTRATIVE

	Adviser Directs and decides	AI Executes within authority
01 SIGNAL	Receives the flag, with the evidence behind it.	Notifies the Morgans' projected income gap widened after an account change.
02 HUMAN DECISION	Meets the Morgans; together they decide to refresh the income analysis and involve their CPA.	Prepares the options, then records the decision and its scope.
03 AUTHORIZED EXECUTION	Can review progress, change authority, or stop the work at any point.	Refreshes the analysis, creates the tasks, and sends the approved request to the verified CPA.
04 VERIFIED RESULT	Reviews the record and resolves the exceptions.	Confirms what completed, what failed, and what still needs a decision.

Dark cells mark who leads each step. Anything outside the authorized scope returns to the adviser.

This scenario is illustrative. It makes the boundaries visible: the decision, the information needed, the permitted actions, and the point where the system must return to a person. CFP Board's technology guidance reinforces the professional's responsibility to exercise judgment when evaluating AI work. Greater automation does not remove that responsibility. [3]

IN CLOSING

Advisers direct. AI executes.

Advisers direct

Set intent

Make decisions

Resolve exceptions

AI executes

Carries authorized work

Verifies the result

Records every step

The opportunity is not a faster version of yesterday's process. It is proactive care: the whole household in view, drift caught early, and authorized work carried through to a verified result, with the adviser directing and deciding.

This is the thinking behind Helmsted. Helmsted is built as the harness: the household record, the permissions a firm grants, the adviser's decision points, and the record of what was done live in one platform. Advisers set intent and resolve exceptions; AI carries out the authorized work inside the boundaries the firm defines, with the evidence visible.

Our security program pairs continuous monitoring through Vanta with an independent SOC 2 Type I examination by Sensiba LLP, which issued a clean report covering Security and Confidentiality, available under NDA.

Build the harness now. Earn the trust that delegation will require.

See how Helmsted brings the household into view.

Thirty minutes with the team, and every security question answered.

[Book a demo · helmsted.ai/demo](https://helmsted.ai/demo)

Sources and further reading

Research checked September 24, 2026. Bracketed numbers in the text refer to the sources below.

1. **NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1 (July 2024)**. Section 2.2 discusses confabulation and next-token prediction; section 2.9, Information Security, discusses direct and indirect prompt injection; the profile also offers voluntary risk-management actions. It does not establish a Reg S-P compliance certification. [Read the profile](#).
2. **Nouha Dziri et al., Faith and Fate: Limits of Transformers on Compositionality, NeurIPS 2023**. Research on multi-step reasoning in language models; the abstract notes that autoregressive performance can rapidly decay as task complexity increases. [Read the paper](#).
3. **CFP Board, Guide to CFP Board's Technology Standard (2024)**. Relevant passages: printed pp. 2–5 on AI, professional judgment, and the technology platform; pp. 7–10 on diligence and evaluation of outcomes. Professional guidance for CFP® professionals; not an SEC rule. [Read the guide](#).
4. **CFP Board, Generative AI Ethics Guide: A Checklist for Upholding the Code and Standards (February 25, 2025)**. Additional reading for CFP® professionals considering AI use. [View the publication and checklist link](#).
5. **Anthropic, Building Effective Agents (December 19, 2024)**. Engineering guidance on agent design, including ground truth from the environment at each step, checkpoints for human feedback, and stopping conditions. Industry guidance, not a standard. [Read the article](#).
6. **Anthropic, Measuring AI Agent Autonomy in Practice (February 18, 2026)**. A study of how people oversee a coding agent in practice, including auto-approve and interruption rates by user experience. Software-development context; not a study of financial advice, and not evidence that graphical interfaces will disappear. [Read the study](#).

Also in The RIA's Chartbook. No. 01, *Client Data and the Law*: what Reg S-P expects as AI enters the workflow. No. 02, *Every Vendor Is Part of Your Safeguards*: evaluating the technology that holds client information. No. 03, *Faster Check-ins, or Better Health?*: the harness behind proactive, AI-enabled advice.

Educational information, not legal advice. Applicability and incident decisions depend on the firm's circumstances and other relevant requirements. Workflow examples describe an operating vision, not a representation that all features are currently available in Helmsted. The medical-practice, rolling-ball, spreadsheet, and autopilot comparisons, and the Morgan household, are illustrations, not statements about any healthcare provider, AI system, or aviation practice. Images are AI-generated editorial illustrations; people and settings are fictional.

Helmsted

One secure, AI-native platform for the independent RIA. Learn more at helmsted.ai, or write to hello@helmsted.ai.