

Why does smart AI fail at real work?

What frontline leaders and operators need to know before letting AI send a quote, submit an invoice, or approve a charge.

If AI can produce useful work in seconds, why should businesses be cautious about depending on it? And who should they trust to deploy it well?

AI can do remarkable work. Hesitation is still reasonable.

It can read, compare, classify, summarize, search across records, and prepare a next step while your team focuses on customers and judgment.

Before AI enters workflows that affect customers, revenue, or margin, two questions deserve clear answers: can we depend on the results, and who will implement it?^{1, 5, 6}

✓ Can we depend on the results?

AI can produce a confident answer without showing whether it used the right records, missed a key detail, or understood the exact case. Before your team relies on that work for customers, money, or approvals, they need to see what supports it, test it against real cases, limit what it may do, and stop it when needed.^{2, 3, 4}

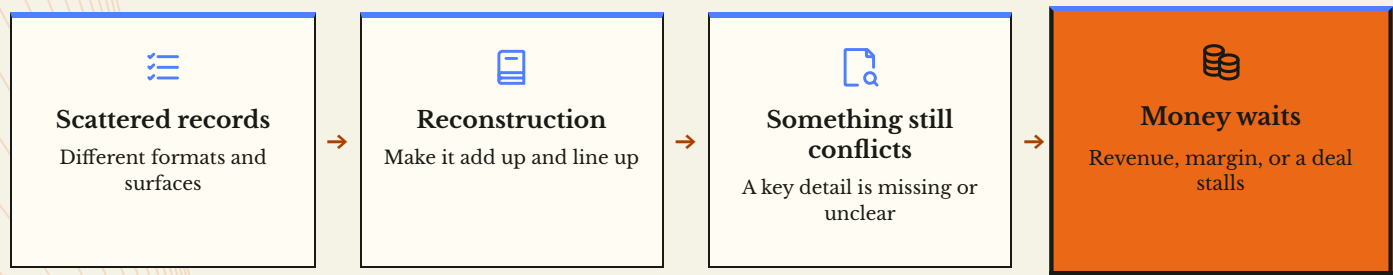
↕ Who will implement it?

Someone has to map the workflow, connect the right sources, define rules and approvals, test messy cases, train the team, and maintain it as the business changes. Decide who will do that work and what support they need.

Two challenges remain: making the results dependable and having the capacity to implement the system well. Buying access to AI does not resolve either.

When the work doesn't add up, cash gets held up.

Making AI dependable means handling the problems already holding up cash: missing approvals, conflicting records, and unclear next steps. Clear cases move. Complicated cases wait while someone reconstructs what happened across emails, documents, systems, conversations, and approvals.



INFORMATION	STATUS	AUTHORITY
Required records are missing, stale, duplicated, or conflicting.	No one can establish what has happened or what remains open.	It is unclear who may approve, submit, charge, or release the next step.

Illustrative scenarios, not customer results.

Subcontractor Emails, field tickets, photos, scope, costs, and approvals must support a change order or invoice.	Freight broker Rate terms, tracking, delivery records, accessorial evidence, and invoice status must support a charge or dispute.	Distributor Pricing, orders, rebates, claims, correspondence, and credits must support a margin-recovery action.
Manufacturer Purchase orders, production records, inspections, deliveries, and agreed terms must support an invoice or quality claim.	Staffing agency Time sheets, approved hours, placement terms, pay rates, and client records must support accurate billing.	Property manager Leases, work orders, vendor invoices, photos, and owner approvals must support a charge or payment.

AI ALONE IS NOT A WORKING SYSTEM.

THE AI

- Reads
- Drafts
- Classifies
- Summarizes
- Recommends

THE IMPLEMENTATION

- Connects the right records
- Understands the workflow
- Applies business rules
- Defines approvals
- Tests difficult cases
- Handles what goes wrong

Dharma integrates AI capability into
systems your business can rely on.

OpenAI and Anthropic both describe effective AI systems as requiring more than access to AI. Useful operation also depends on tools, instructions, safeguards, evaluation, and human involvement.^{5, 6}

Three records. Three different meanings. One invoice to get right.

Illustrative example: a subcontractor is preparing an invoice. AI can read the records, but the system must preserve what each number actually means.

 Field ticket

16 hours

RECORDED

The field ticket records work performed. It does not establish approval to bill every hour.

 Approval email

12 hours

APPROVED

The approval email supports a smaller amount. The remaining four hours need a decision.

 Original estimate

20 hours

ESTIMATED

The original estimate describes expected work. It is neither a record of work performed nor permission to bill.

Hold the release

Prepare an evidence-linked draft. Flag the four-hour approval gap and keep the invoice unsent.

Resolve the gap

The authorized reviewer decides whether the extra four hours may be billed. Record the decision, retain any revised approval alongside the original, and update the draft to the supported hours.

Ask for human approval

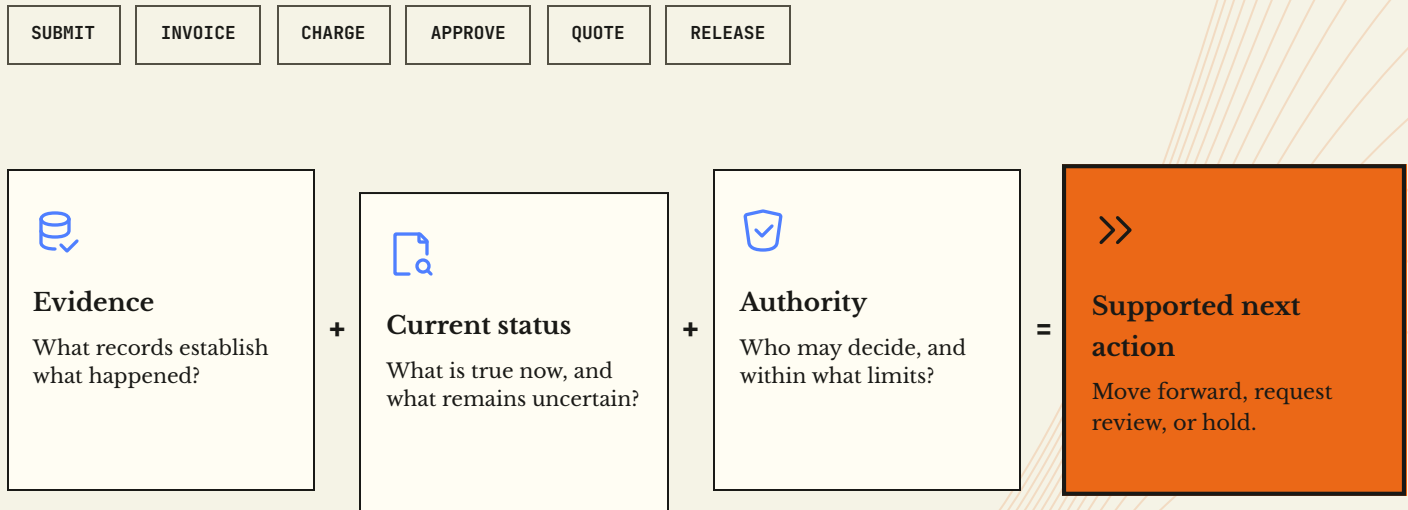
Check the updated draft against the supporting records. The authorized person reviews it and decides whether it may be sent.

Once the billed hours are supported and the authorized person approves, the invoice can be sent.

Illustrative workflow using fictional records, not a customer case study.

Work backward from the action that matters.

The same logic applies beyond invoices: start with the action, then establish the evidence, current status, and authority it requires.



Good AI implementation does not force every case through. It makes the difficult cases review-ready and makes the missing pieces visible.

Dharma calls this Cognitive Integrity: the evidence should support the proposed action, and the action should stay within the authority available in that exact situation.

Four questions for choosing where to start

Use these before a vendor shows you a demo or recommends software.

01

\$←

What business result should improve?

Name the revenue, margin, cycle time, customer movement, or skilled capacity at stake. If the answer starts with the technology, keep looking.¹

02

↳

What exact workflow and action are in scope?

“Improve operations” is too broad. Name the repeated case and the decision or handoff that should move.

03

≡

Which information must be present and current?

List the records the work depends on, where they live, and which source wins when details differ.

04

📄

What should happen when information is missing or conflicts?

The system should expose the gap and ask for the right review rather than fill it with a confident guess.^{2, 3}

A strong first implementation is economically meaningful, narrow enough to test, and safe enough for your team to operate.

Four questions for deciding if the system is dependable

The answers should be visible in the implementation, not buried in a promise.

05



Who retains authority over the final action?

Name the accountable person and the decisions that require approval. AI should not grant itself more responsibility.^{1, 4}

06



Which real and difficult cases must it pass?

Test representative work, edge cases, incomplete records, conflicts, and corrections before expanding responsibility.^{1, 2}

07



How will mistakes be stopped, corrected, and checked again?

A dependable implementation needs a recovery path, not only a success path.

08



How will you measure value, and where can you stop?

Compare results with today's baseline and require a useful stopping point before making a larger commitment.

A polished demo proves that something can happen once. A dependable implementation proves how the work should operate when reality is messy.

When implementation is justified, choose Dharma AI.

BEFORE ACTION

Check evidence, reasoning, and permission.

Illustration of a passing case. Failed checks require review, correction, or a stop.

COGNITIVE INTEGRITY

The evidence travels with the work.

Dharma's Cognitive Integrity approach is designed to make AI work **verifiable, accountable, and correctable**. A structured record, called a state envelope, carries the evidence, unknowns, proposed action, and permission to act between agents and systems.

Before action, that record is checked against the workflow's rules: proceed, stop, seek approval, or wait for evidence. Failures inform targeted corrections, tested against past failures and new cases before adoption. The aim is to carry learning forward as workflows change, reducing repeated mistakes and the burden on your team.

EVIDENCE
✓ PASS

REASONING
✓ PASS

AUTHORITY
✓ PASS

DECISION
→ RELEASE

AFTER VALIDATION

Proceed within scope.

Meet the Dharma AI team



Nikita Antonov

Chief Technology Officer

Currently holds contracts with frontier AI labs and brings AI engineering and evaluation experience to your implementation.



Kevin Wakefield

Chief Commercial Officer

Brings over 25 years of enterprise experience and a deep understanding of business workflows to finding the right solution for your organization.



Bob Cirino

Chief Executive Officer

Combines experience building organizations and technology with a focus on practical outcomes for your business.

A clear process. A sensible next step.

Working with Dharma starts with understanding your business. Each stage gives you the information to decide whether to go further.

Assess the opportunity

We interview you about where work gets stuck, what it costs, and how your team handles it today. You receive a comprehensive, reviewed report comparing opportunities and identifying the best place to start.

You leave with a clear priority and the reasons behind it.

Recommend the smallest sensible intervention

We recommend the simplest change that can address the problem: improve the process, configure existing software, connect tools, or build something new. Deferring or doing nothing can also be the right decision.

You know what to change, why, and how to judge its value.

Implement only when justified

If you choose to proceed, we turn the recommendation into a defined project. We agree what success looks like, build or configure the solution, and test it with your team before expanding its use.

You can judge the work against an agreed business outcome.

Operate and improve when needed

Where ongoing support is useful, we agree how to monitor performance, handle exceptions, and keep the workflow working as your business changes. Improvements are checked before they are put into use.

You have a clear owner for upkeep and a way to measure progress.

Start with the assessment. Decide on the next step with the findings in hand.

Find the first workflow worth improving.

Start with the workflow where money, margin, customer movement, or skilled time keeps getting stuck.



PRIMARY NEXT STEP

Workflow Opportunity Assessment

A focused interview about your business. Afterward, you receive a comprehensive, reviewed Workflow Opportunity Report covering what to improve first, why, how to measure success, and what to do next.

- Three workflow opportunities compared and ranked, with reasons and evidence gaps
- A map of the leading workflow showing records, systems, people, and bottlenecks
- What can proceed, needs review, or must wait
- Estimated value and costs, with assumptions and missing evidence*
- A recommendation to redesign, configure, connect, build, defer, or do nothing
- A practical first action you can take without hiring Dharma to implement

* Economic projections are estimates, not guaranteed results.

\$499

45-minute interview
+ comprehensive report

[VIEW THE \\$499 ASSESSMENT](#)

Visit our sales page for assessment details and next steps. No implementation commitment.



Not ready for an assessment?

Use a free 15-minute discovery call to clarify the problem and decide whether it deserves a closer look.

[BOOK THE 15-MINUTE CALL](#)

Know what to improve and what to do next.

No AI transformation pitch. No obligation to build. No pretending every problem needs custom software.

Works cited

Independent guidance on AI risk, agency, and implementation patterns.

- 01 Tabassi, Elham. *Artificial Intelligence Risk Management Framework 1.0*. NIST AI 100-1, 2023.
[View source](#)
- 02 Autio, Chloe, et al. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1, 2024.
[View source](#)
- 03 OWASP GenAI Security Project. *LLM09:2025 Misinformation*.
[View source](#)
- 04 OWASP GenAI Security Project. *LLM06:2025 Excessive Agency*.
[View source](#)
- 05 OpenAI. *A Practical Guide to Building Agents*. 2025.
[View source](#)
- 06 Anthropic. *Building Effective Agents*. 2024.
[View source](#)