


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?

When the work doesn't add up, cash gets held up. Could AI help?

Missing approvals, conflicting records, and unclear next steps hold up cash. Clear cases move. Complicated cases wait while someone reconstructs what happened across emails, documents, systems, conversations, and approvals.



Use cases:

Subcontractor A change order or invoice must be supported by emails, field tickets, photos, scope, costs, and approvals.	Freight broker A charge or dispute must be supported by rate terms, tracking, delivery records, accessorial evidence, and invoice status.	Distributor A margin-recovery action must be supported by pricing, orders, rebates, claims, correspondence, and credits.
Manufacturer An invoice or quality claim must be supported by purchase orders, production records, inspections, deliveries, and agreed terms.	Staffing agency Accurate billing must be supported by time sheets, approved hours, placement terms, pay rates, and client records.	Property manager A charge or payment must be supported by leases, work orders, vendor invoices, photos, and owner approvals.

Feed all of your information to an AI model and ask, “What needs to happen next?” Without the right rules, checks, and approval steps, you risk a beautifully worded answer that is **confidently wrong**.

That’s not an AI system your business can rely on.

AI can do remarkable work. Hesitation to rely on it is reasonable.

AI can take work off your team's plate. But can you depend on it, and who will make it work for your business?^{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.

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 AI 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} Dharma AI requires agents to prove Cognitive Integrity before taking a consequential action.

WHAT GOOD AI IMPLEMENTATION REQUIRES

Establish the facts. Operate within limits. Execute logical work.

Before an AI system takes a consequential action, like sending an invoice, submitting a claim, or releasing a payment, it must establish what that action requires and show that those requirements are met. Good implementation builds three checks into the workflow.

EVIDENCE

What does each record establish?

Find the relevant records, verify their sources, and understand their context. Distinguish estimates, completed work, and approvals. Organize those facts into fields the system can check against the specific requirements of the output.

Keep the meaning, source, and current status attached to each fact.

PERMISSION

What can proceed?

Apply business rules and approval limits to the evidence. Determine what may proceed, what can only be drafted, and what requires permission. A record that work happened does not, by itself, authorize billing it.

Evidence supports the action. Rules and permissions govern it.

RESOLUTION

What will close the gap?

Identify what is missing or conflicting and where clarification can come from. Let the agent retrieve records or request clarification within its permissions. Escalate when human judgment or approval is needed.

Check each new answer against the requirements before proceeding.

↳ New information goes back through the evidence and permission checks.

This is how you move from an AI model's opinion to a calculated outcome you can verify and rely on.

EXAMPLE: AI IMPLEMENTED FOR SUBCONTRACTOR BILLING

The crew's been paid. You haven't. And you can't find approval.

The job started at 12 hours. A change order added four. Your crew did all 16. You remember a go-ahead, but can't find approval to bill the extra work. Did you miss an email, or was approval never sent?

BEFORE

You search email threads to match the extra work to the change order and struggle to find approval. Do you bill the original 12 hours and eat the four you couldn't verify, or bill 16 without support and risk a dispute? The next job's already underway.

AFTER

The AI workflow checks your connected communication and record systems. It finds approval for 12 hours, but none for the extra four in the records checked. It identifies who can approve them and drafts the request with the change order and field ticket attached. You review and send.

ORIGINAL APPROVAL

12

 hours

You have approval to bill the original 12 hours of work.

CHANGE ORDER

+4

 hours

Four more hours of work are requested on the same job. In this example, approval to bill them comes separately.

FIELD TICKET

16

 hours

The crew records the original work and the extra four hours as completed.

APPROVAL NOT FOUND

+4

 hours

No matching approval was found in the records checked. These hours need approval before billing.

Identify the gap

Link the completed work to the change order and original approval. Show which records were checked and any sources that could not be accessed. Keep the four unapproved hours flagged.

Draft the request

Prepare a message to the person who can approve the extra hours, with the supporting records attached. You review and send it; the request stays open until resolved.

Update the invoice

When approval arrives, verify the approver, job, and four extra hours. Then update the draft to 16 hours with both approvals attached. You decide when to send the invoice.

Draft ready for your review: “Please confirm approval to bill the four additional hours on the attached change order. The field ticket shows 16 hours completed, including the 12 already approved.”

Is this worth implementing?

Start with the business case. These four questions help you decide what deserves attention, how small to start, and whether AI is the right approach.

01 What recurring problem is costing the business?

Identify the delayed cash, lost margin, repeated effort, or customer friction. Establish how often it happens and what it costs today.¹

02 What specific work should the system take on?

Define the starting event, intended outcome, and boundaries. Make the first implementation small enough to evaluate before expanding it.

03 Could a simpler change solve the problem?

Consider improving the process, configuring existing software, or connecting systems before building an AI workflow.

04 Would the improvement justify the full cost?

Compare today's baseline with expected value, including implementation, operation, and human review. Define a useful stopping point before making a larger commitment.

You do not need every answer in advance. An assessment should help establish whether the opportunity is worth pursuing.

If the business case holds, examine the proposed implementation.

Can you trust the proposed implementation?

Ask your implementation partner for concrete answers: the evidence, permissions, tests, and ongoing ownership behind the proposal.

05 How will it establish that an action is justified?

Show how records are found and interpreted, checked against the action's requirements, and retained as supporting evidence. Missing or conflicting support must stay visible.^{2, 3}

06 What can it do on its own, and when must it ask?

Define permissions, what happens when evidence is missing, and who makes decisions the agent cannot. Responsibility must be granted by an authorized person.^{1, 4}

07 How will you test it and handle failures?

Demonstrate ordinary and difficult cases, how mistakes are contained, and how corrections are checked before work resumes.^{1, 2}

08 Who keeps it working after launch?

Name the owner for monitoring results, handling exceptions, and adapting to changes. Agree how to decide whether to continue, adjust, or stop.

When the business case holds and the proposal shows a dependable way to deliver it, AI implementation is justified.

Choose a partner who can make those answers concrete in your business.

When implementation is justified, choose Dharma AI.

BEFORE ACTION

Check evidence, reasoning, and permission.

EVIDENCE

✓ PASS

REASONING

✓ PASS

AUTHORITY

✓ PASS

DECISION

→ RELEASE

AFTER VALIDATION

Proceed within scope.

COGNITIVE INTEGRITY

Dharma AI makes agents earn their right to act.

Our team has been deploying AI systems for more than eight years. We've seen the trouble capable AI models can create when poorly implemented. That experience led us to build Dharma AI and help businesses implement AI worth using.

Dharma AI'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.

Before action, that record is checked against the workflow's rules: pass, block, escalate, withhold pending additional evidence. Failures inform targeted corrections, tested against past failures and new cases before adoption so the system can adapt as workflows change.

Every Dharma AI implementation includes the platform we built for reliable AI.

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.

COGNITIVE INTEGRITY IN PRACTICE

Dharma AI makes AI agents earn the right to act.

Connecting AI to a workflow makes action possible. Dharma AI adds the evidence, authority, and confirmation needed to make consequential action accountable.

IMPLEMENTATION ASKS

Can it complete the task?

DHARMA AI ALSO ASKS

Should it act, and can it prove the result?

BEFORE ACTION

Make the proposed decision inspectable.

A state envelope carries the action, supporting evidence, remaining unknowns, and claimed authority.

The agent cannot hide a gap inside a confident answer.

AT THE DECISION

Check whether the action is earned.

Dharma AI's integrity contract tests perception, reasoning, and authority against the workflow's rules. The result is explicit: release, block, escalate, or wait for evidence.

AFTER ACTION

Confirm what actually happened.

Permission to act is not proof of completion. The receiving system must confirm the result before the agent can claim it is done. Failures produce a precise correction and a new test.

WHY IT MATTERS

Dharma AI delivers the workflow and the conditions that govern when it may act: clearer accountability, more useful autonomy, and less repeated cleanup.

How to implement AI with Dharma AI.

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

 Assess FIND WHERE WORK GETS STUCK	A clear priority. 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.
 Recommend THE SMALLEST USEFUL CHANGE	A practical recommendation. Choose the simplest change that addresses the problem: improve the process, configure existing software, connect tools, or build. You know what to change, why, and how to judge its value. Deferring can also be the right decision.
 Implement ONLY WHEN JUSTIFIED	A tested workflow. If you choose to proceed, we agree the scope and what success looks like. We build or configure the solution and test it with your team before expanding its use. You judge the work against an agreed business outcome.
 Improve ONGOING SUPPORT WHEN NEEDED	Clear ownership as things change. Where ongoing support is useful, we agree who monitors results, handles exceptions, and keeps the workflow working as your business changes. Improvements are checked before use, with a clear way to measure progress.

Start with the assessment.

Decide what comes next with the findings in hand.

Let's identify 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 AI 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](#)