Ontollo.

WHITE PAPER · SEPTEMBER 2026

The Operations Leader's Guide to AI in Industries that Build and Run the World

Why AI stalls in construction, infrastructure, and energy and what is needed to fix it.

Contents

---

1. Executive summary

Most asset-intensive organizations have made serious investments in systems of record. Enterprise resource planning, project controls, procurement, scheduling, payroll. Each system is authoritative for its own facts, and each was expensive and difficult to adopt. None of that investment was wasted.

What no system holds is the set of relationships between them: which commitment belongs to which job, worked by which crew, against which scheduled activity, drawing on which asset. That connective tissue exists today only in the judgment of experienced people and in spreadsheets rebuilt each month. It is unwritten, inconsistent between divisions, and often unreadable
by software.

That missing layer has a formal name: an ontology. An ontology is the agreed definition of the

things important to an organization, its business, and how they relate to one another. It is the reason enterprise AI pilots keep producing confident answers to partial questions, and it is the single decision that determines the ceiling on what AI can eventually do for an operation, as well as the floor on what that AI costs to run.

As companies give employees access to AI tools and training, they all reach the same conclusion: there are limited productivity gains for individuals unless systems can be connected together. Effective use of ontologies with AI makes that happen and can reduce token counts by 89%.

This paper covers four things:

- Why systems of record, working exactly as designed, cannot tell an operator when their picture has gone stale
- What an ontology is in plain terms, with a worked example from construction
- Why the ownership of that model is an executive decision rather than a technical one
- Five questions that separate credible options during an evaluation

2. AI gains for physical operators lag behind

The productivity gains of the past several years have landed almost entirely with people who work in front of screens. Drafting, summarizing, searching, coding. Those gains are real and they arrived quickly.

The organizations running sites, plants, grids, and fleets are still waiting. This is not because the models are weaker in physical industries. It is because AI needs a current picture of what is actually happening, and for most physical operations that picture does not exist in a form software can reach.

An operator can describe the state of a job in a two minute conversation. No system in that operator's organization can assemble the same answer, because the answer depends on five systems that were never designed to speak to each other, plus what the superintendent saw this morning.

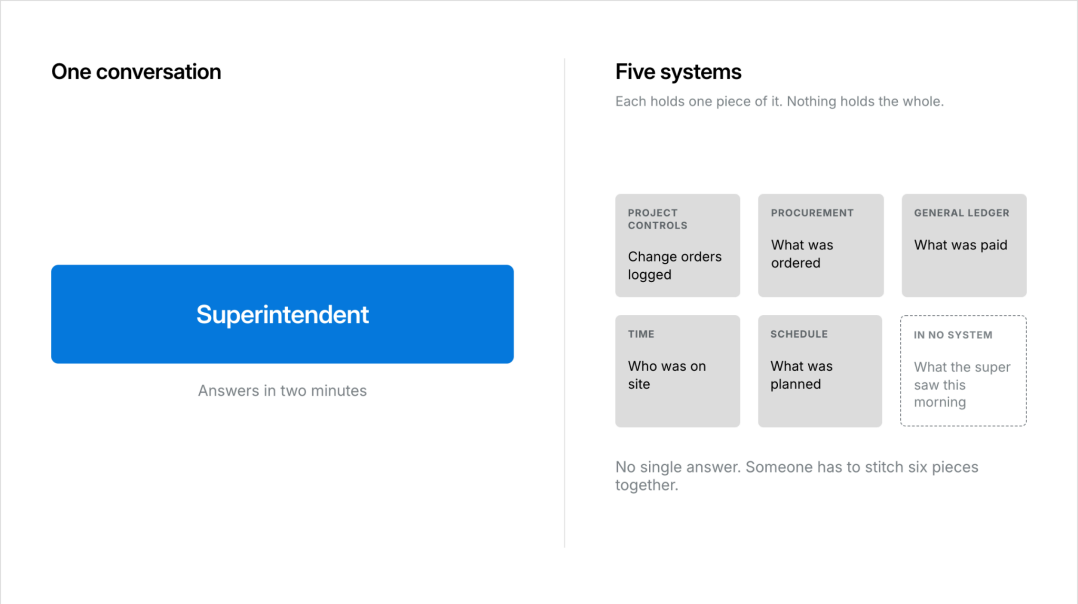

3. What a system of record does well, and what it structurally cannot do

Systems of record such as CRMs and ERPs answer one question extremely well: what happened. Who submitted the request, when it was approved, which revision is current, who signed off. On a project generating tens of thousands of documents, answering that reliably is the difference between a defensible position and a claim that gets lost. That capability was hard to build and it took the industry a long time to adopt. It is also, structurally, a description of the past.

A record is only as current as the last person who updated it, and it has no way of knowing when that was long enough ago to matter. A project board that stopped being updated three weeks ago looks exactly as confident on day twenty-one as it did on day one. The data is clean. The system is working as designed, but the picture doesn’t reflect the current reality. In addition, nothing in the architecture can flag it, because a record stores what it was told rather than whether what it was told is still true.

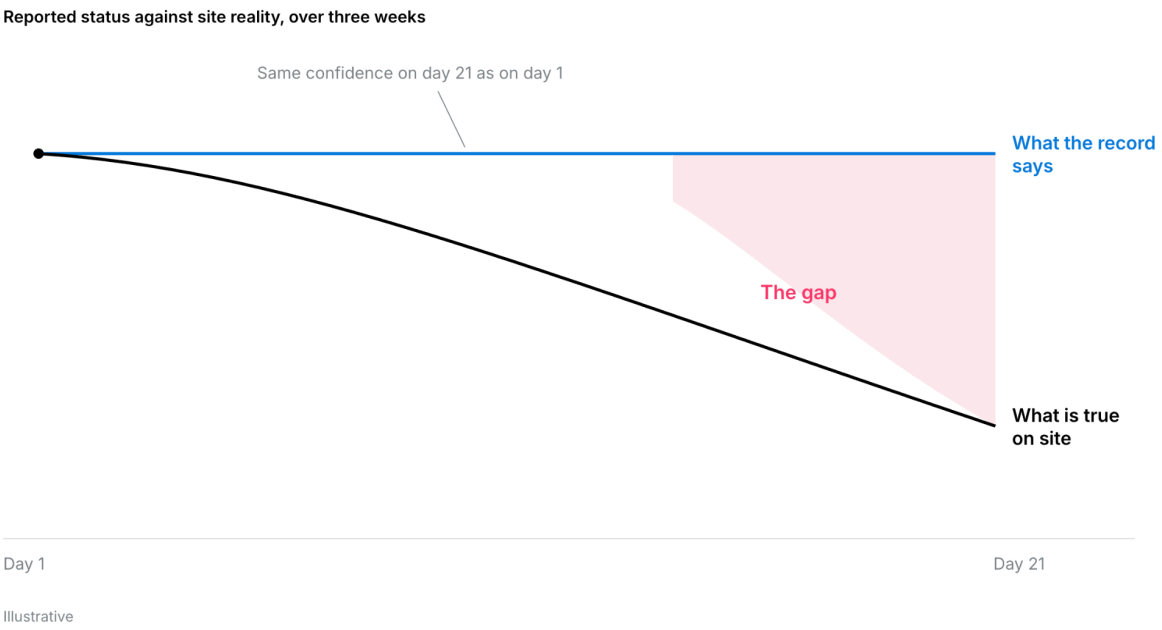

Consider the scale at which this operates. The Navigant Construction Forum studied 1,362 projects containing more than a million requests for information and found that more than one in five never received an answer at all. That research dates to 2013 and the tooling has improved considerably since. The shape of the finding has not. Every one of those items sat in a system that knew it existed, with a date, a recipient, and a status. Each record was complete and correct by its own standards. What none of them could do was signal that an open item was holding up a crew, or that the system's own picture of the project had drifted from what was happening on site.

The gap between the plan and the field is not a data quality problem. It is a timing problem, and it hides in plain view.

4. Why adding AI to that picture does not work

The instinct is to point an AI agent at the record and let it find what people are missing. That produces something genuinely useful. An agent can search, summarize, and draft against what a record contains, and it saves real hours.

With that said, the AI agent also inherits the record's blind spot. An agent will act confidently on information that stopped being true two weeks ago, and it will not hesitate, because nothing in its inputs signals doubt. Automating action on a stale picture does not fix the picture. It arrives at the wrong place faster.

The evidence at industry scale is consistent with this. MIT's Project NANDA reviewed hundreds of enterprise deployments and found that 95% of generative AI pilots produced no measurable P&L return against an estimated $30 to $40 billion of spend. The report is direct about the

cause, and it is not model quality. It is integration and context.

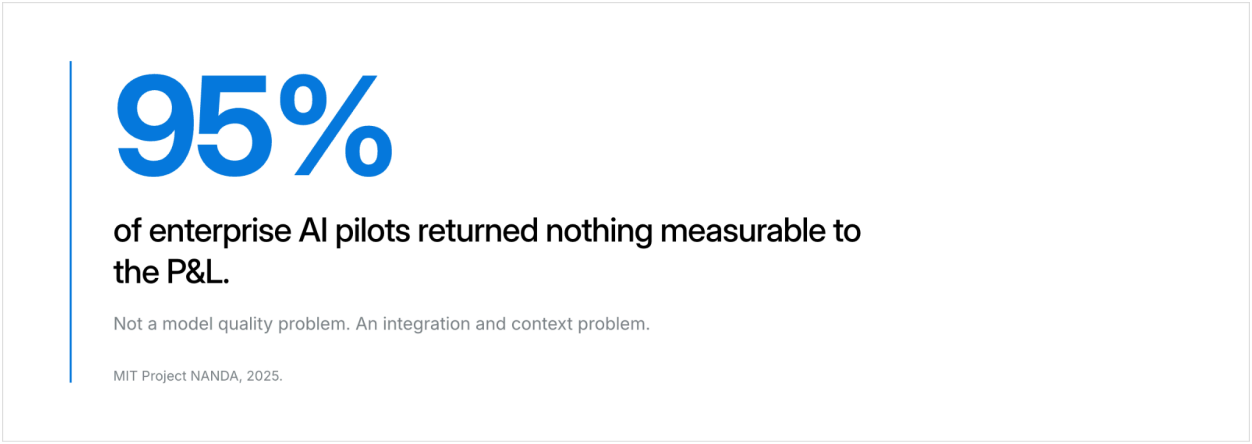

.

Operations leaders tend to recognize this before analysts do. An assistant asked to produce a job health report reaches whichever slice of data it has access to. That slice rarely includes commitments, invoices, schedule, and ledger at the same moment. So the assistant returns a confident answer to a partial question, which is more dangerous than no answer, because someone acts on it.

5. Ontology: What it is and why it matters

This is where an ontology comes in. An ontology is an agreed definition of the things an organization runs on and how they connect to each other. Ontologies are commonly referred to as the “context layer” that gives AI the information it needs to operate effectively, respond accurately, and stay within the guardrails set by an organization.

Ontologies consist of two parts. The first is the entity: what counts as a job, a commitment, a crew, an asset, a delay. The second is the relationship: which of those things depends on which others, and what a change in one implies for the rest.

Software has historically been built around the first half only. Each platform defines its own objects well and connects to its neighbors through fixed integrations, one process at a time. Facts therefore arrive at their destination with the context stripped off. An operational knowledge graph inverts that arrangement. It stores entities and the relationships between them, which is much closer to how the work is actually organized and how experienced people already think about it.

Fixed integrations

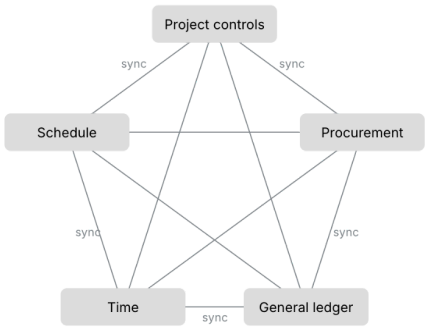

Facts arrive with the context stripped off

Operational knowledge graph

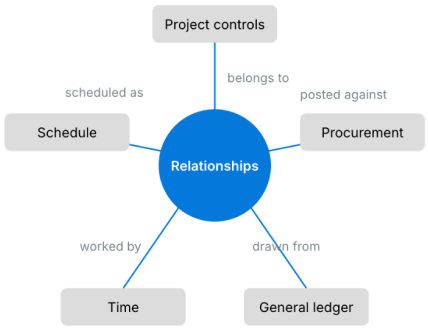

Facts arrive with what they mean attached

6. An example in construction

Let’s apply this to an example that is common in the construction industry.

A field team places a material order to keep a pour on schedule. Nothing about that decision is unusual. Now trace what each system knows.

- The **project controls system** knows a budget line exists for that scope, and recalculates a forecast whenever anything touches it.
- The **procurement system** eventually knows an invoice arrived, once a coordinator assigns it.
- The **general ledger** knows what has been posted, which is a different number again.
- The **time system** knows crew hours, attributed to a job rather than to the activity the crew was actually performing.
- The **schedule** knows the activity slipped, but not why.

Each system is correct. Each is authoritative for its own facts. No system holds the sentence that matters to an executive: *this invoice belongs to that commitment, under that cost code, on that job, worked by that crew, against that activity, and here is what it does to the number we bid.*

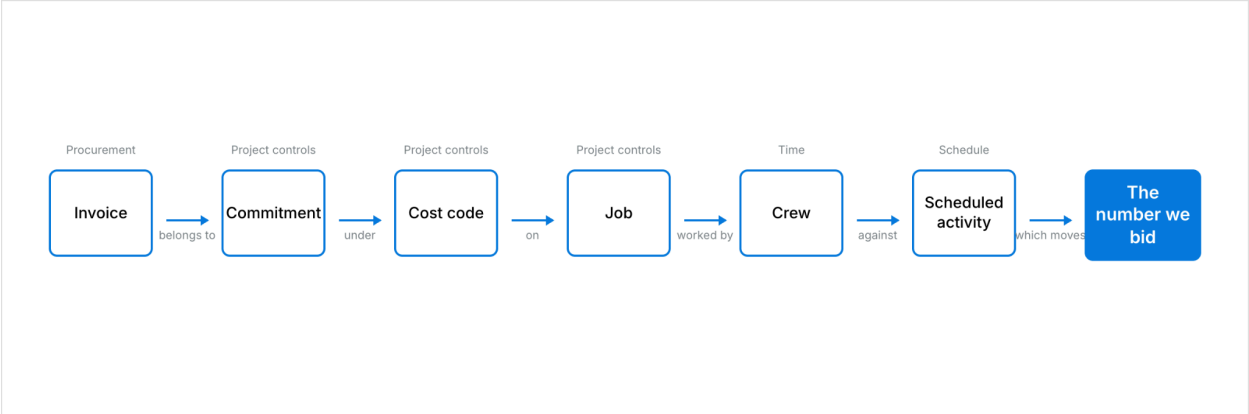

Today a person assembles that sentence by hand, usually at month end, usually under time pressure. By then the decision that created the exposure is weeks old and the opportunity to influence it has passed. Multiply that across a portfolio and the result is an organization that manages from a lagging position, spending most of its effort locating problems rather than resolving them.

An ontology is what allows that sentence to be assembled automatically, in a form that a person or a machine can act on the same day.

The same structure applies outside construction. In energy and infrastructure, the nouns change to work orders, outages, permits, and assets, and the failure is identical: every system holds a legitimate version of the truth, and none of them holds the relationships between the versions.

7. What an ontology gives an operation

Getting the model right has immediate impact on the business:

Answers are assembled rather than retrieved. A number is built from connected facts rather than pulled from whichever export happened to be nearby. Accuracy stops depending on who ran the report and what version of the data they used.

Implications become traceable. The system can answer what a delay on one item means for the rest of the work, not merely that the delay occurred. That is the difference between a notification and a decision.

Every fact carries its source. Any figure can be traced back to the system that produced it. This is the only durable answer to the question every executive asks of a new number, which is how do we know that.

The model extends rather than getting rebuilt. New sources and new signals attach to an existing structure. And critically, definitions outlive systems: when a procurement platform is replaced, the fact that a commitment belongs to a job does not change. Only the source underneath it does. Organizations that wire systems directly to one another pay for that replacement in integration work every time. Organizations that define meaning above their

systems pay for it in a connector.

8. The economics of tokens

There is a second argument for this layer that is newer, and it is worth putting in front of a finance sponsor. An ontology sets the floor on what AI costs to run. An agent that knows how a piece of work connects to everything around it does not have to reconstruct the organization on every request. One that does not know is paying to re-derive the same context indefinitely.

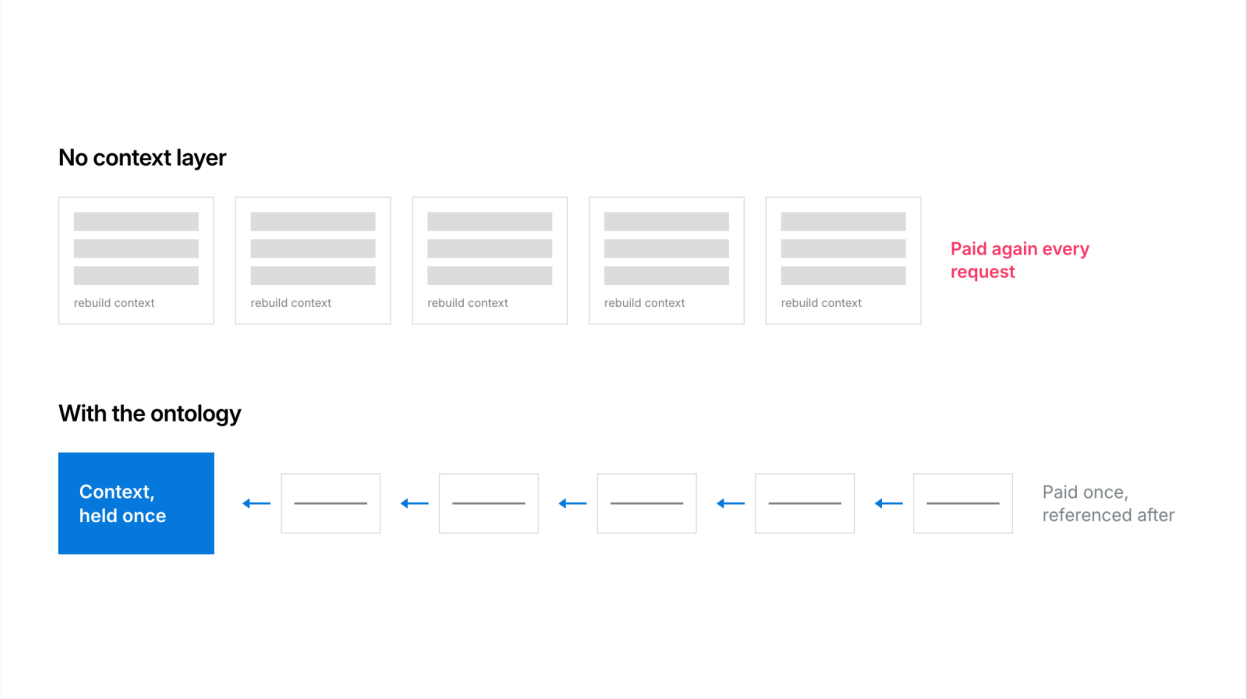

This is not a hypothetical concern. Software companies encountered it first because they shipped AI to millions of users first, and several public examples in 2026 involved features that worked exactly as designed while costing far more to run than anyone had modeled. Physical operations are not exempt from that pattern. They are simply a step behind it.

Most operators will never own a GPU cluster. Control over AI economics has to come from somewhere else, and it comes from a layer that knows which questions are routine and which are load-bearing. Without that layer, the safe default is to send everything to the most expensive available answer, and the cost of that default wipes out any potential productivity gains.

9. Systems of record, reality, and action

It’s important to understand the difference between a system of record, a system of reality, and a system of action. You need all three and you need them to work together to be an AI driven organization for any business.

- A **system of record** tells an organization what happened (e.g., ERP, CRM, project

management tool).
- A **system of reality** tells it what is true right now, and where it is uncertain.
- A **system of action** acts on that, recommending actions to humans or enacting decisions via agents or drones/robots.

The need for all three is important.

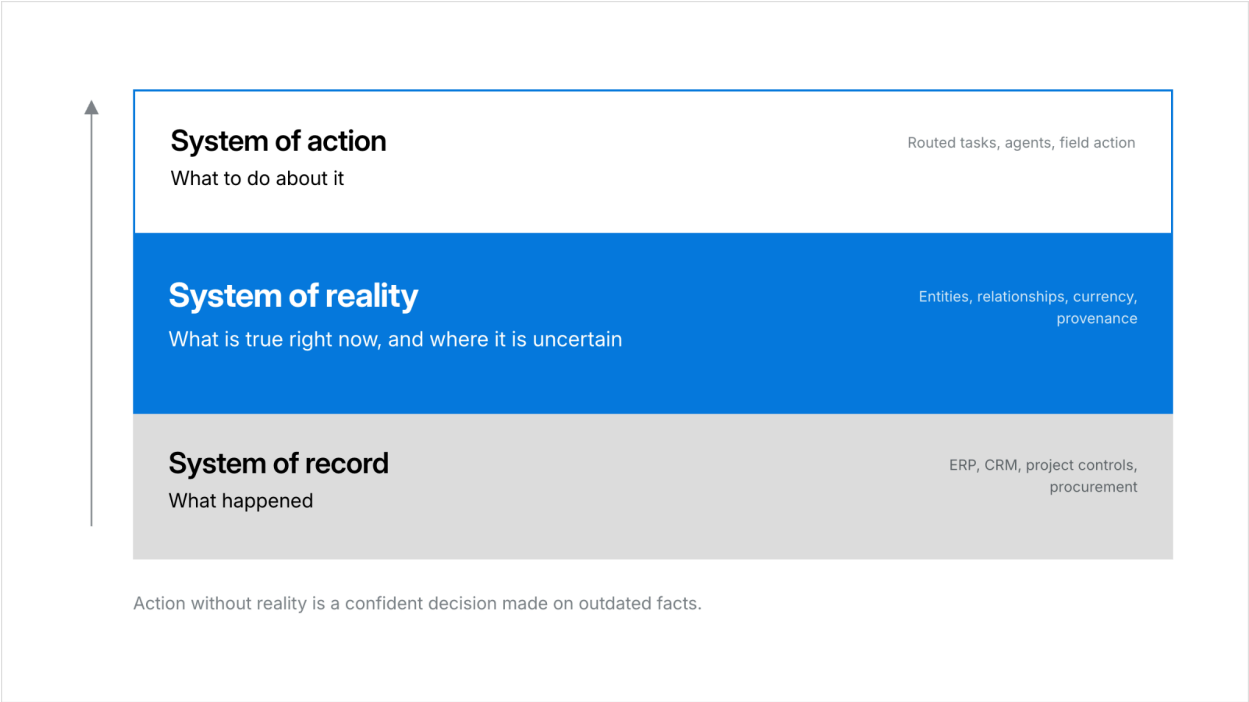

Action without reality is a confident decision based on outdated facts, which on an active site is more expensive than no action at all.

A system of reality can report that the schedule and the field have not agreed in over a week, or that a status has not been touched since the material shipped, or that a section of the work is currently running on assumptions. A system of record cannot do this. It has no representation of its own currency.

10. Why this must be your system, not a product feature

The reasonable counterargument is that the major platforms are building this now, and that eighteen months of patience solves it.

Having built one of these layers before, our view is that this is not primarily a technical problem. A live model of the work requires deciding what the objects are, which relationships are load-bearing, and what a change in one implies for another. Those decisions are specific to how a given organization operates. A self-perform contractor and a construction manager at risk do not share a model of how their work connects, and neither of them has the one that ships in the

box. The same holds between a utility and an infrastructure owner.

It also requires capture where the work happens. A picture of reality that depends on someone returning to a trailer to type is stale before it is entered, which is why voice and scanning on mobile are load-bearing in this design rather than conveniences.

So the question is not whether this model gets built. It is who builds it and who owns it afterward. The organizations that get real value from AI will not be the ones running the most agents. They will be the ones holding a current, connected picture of their own work, and holding it in a form they can take with them.

11. Five questions for any evaluation

Feature lists will not separate the options in this market. Everyone has agents, prebuilt workflows, and orchestration. These five questions will tell you more than any demonstration.

1. **Can it tell us when its own information is out of date?** Push hardest here. The answers are revealing, and a system that cannot say where it is uncertain is asking to be trusted evenly, which no operator does.
2. **Can it tell us what a delay on one thing means for the rest of the work,** or only that the delay occurred?
3. **Can any figure be traced back to the source system that produced it** without a person going to check?
4. **What happens after the finding?** Does it produce an action with a named owner, or another report that someone has to read?
5. **If we change platforms in three years, does the model of how our work connects come with us?** If the model cannot be exported, the organization did not buy software. It rented its own definitions.

12. Where Ontollo fits

Ontollo maintains a deterministic layer beneath an organization's operations that tracks what is actually true in the business rather than a model guessing at it. AI reasons on top of that layer, inside boundaries the customer sets, and every output carries a trail showing how the answer was reached. Nothing changes in a source system without a human confirming it.

Ontollo sits above the systems already in place and replaces none of them, reading project, finance, procurement, payroll, and schedule systems where they are, and capturing what is happening on the job while the work is being performed. What it produces is not another dashboard. It is a task, routed to the right person, with the evidence attached.

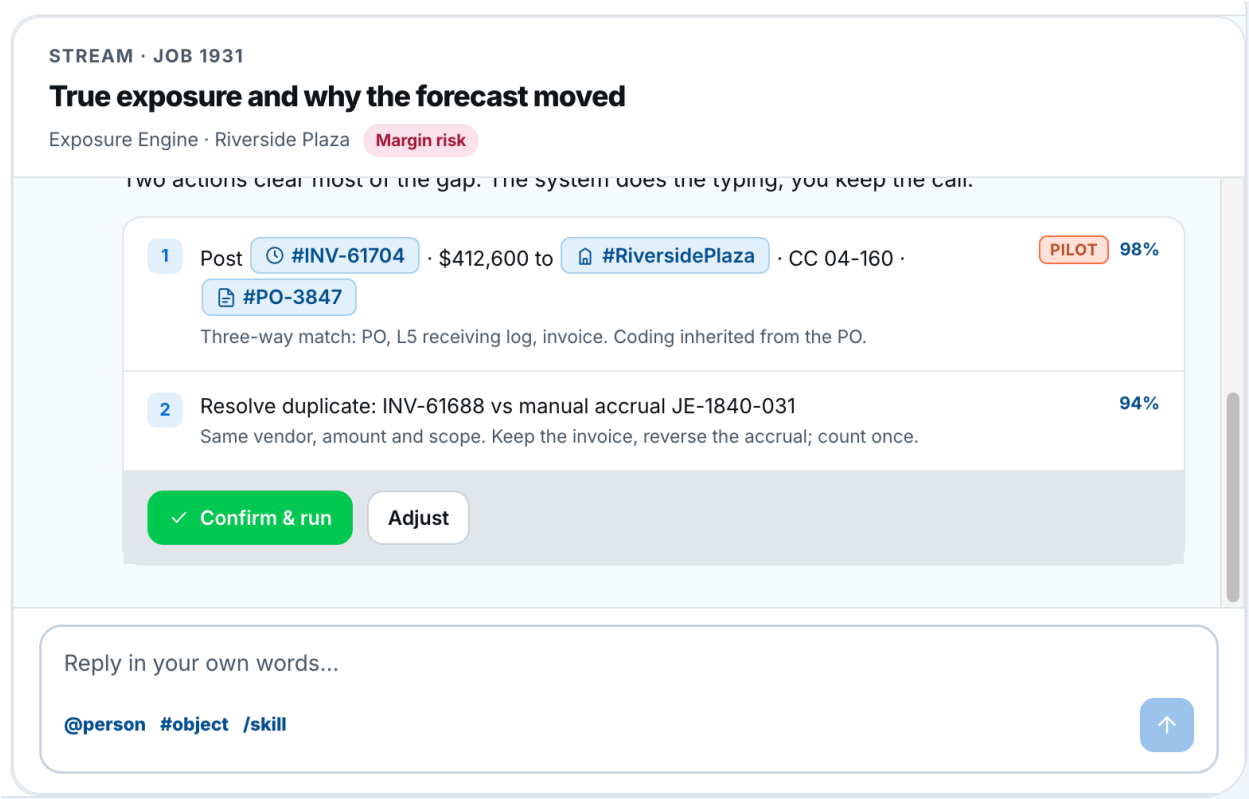

We're helping clients effectively leverage AI in sectors across construction, energy, infrastructure, and healthcare. One of our engagements is with an energy infrastructure platform provider in Australia, scaling past $1B in revenue to meet data center demand across more than fifteen disconnected systems. Their challenge: equipment was being rented externally while identical assets sat idle elsewhere in the group, and leadership knew margin was leaking without being able to locate it. Working with their leadership team, including their CFO, we validated between $8M and $19M per year in avoidable cost. That engagement is now a paid production rollout.

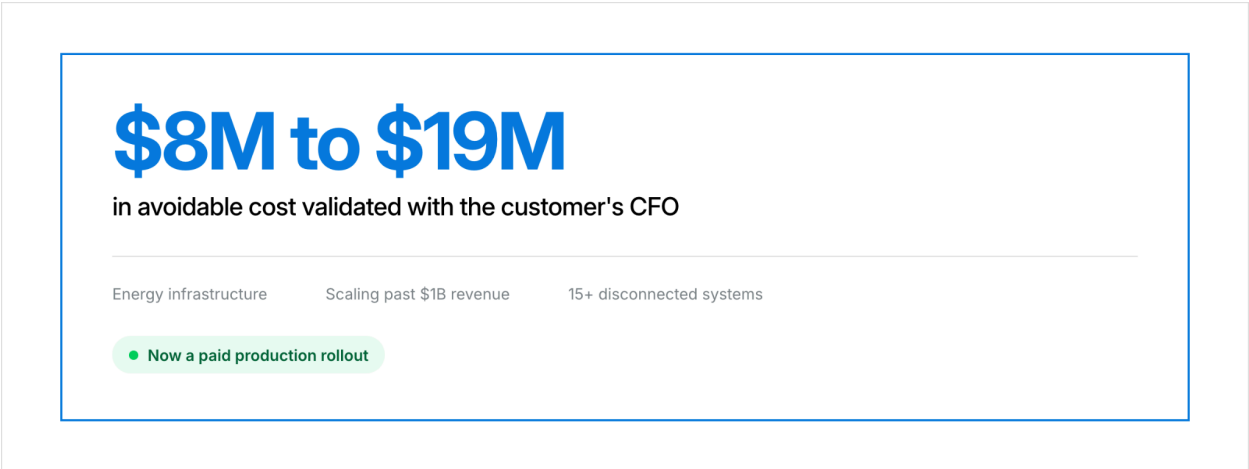

13. Schedule a completely customized demo

If you're interested in learning more about how Ontollo can help you leverage AI effectively,we'd love to connect with you. Our team will give you a demo of our platform, completely customized to fit your exact business needs.

In addition, if there's a fit, we run a 30-day pilot at no cost against your live data: systems mapped in the first week, two or three sources connected in the second, real scenarios tested in the third, and a written findings report in dollars plus a working proof of concept in the fourth, whether or not the engagement continues.

Book a demo and we will walk through your version of this. For readers who want the architecture, ask us for the companion technical white paper, *The System of Execution for Physical Operations*.

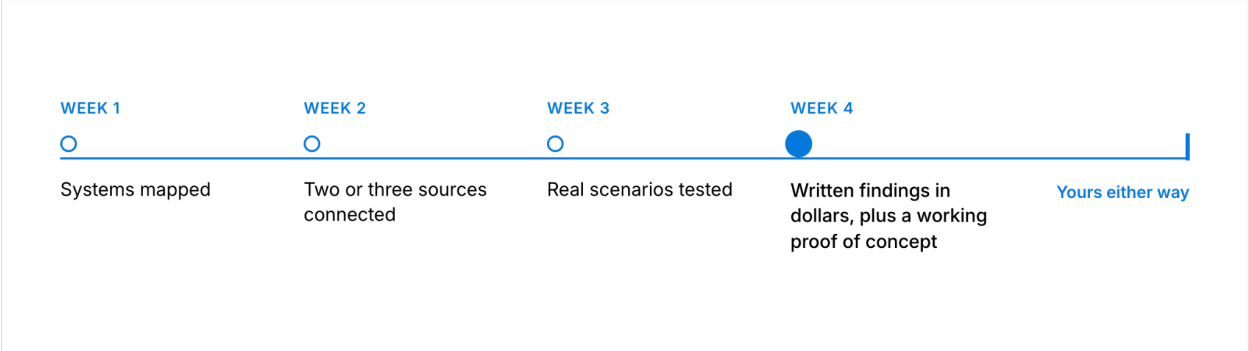

About Ontollo

Ontollo builds the system of reality for physical operations. Founded by Tiffany To and Noah Wasmer, with a combined 40 years building the world's most adopted project management system at Atlassian, large-scale data systems at VMware, and mobile experiences at Apple.

Learn more about Ontollo at www.ontollo.com.