Hellomatik

How we build a *company's ontology*

Design and writing: Isaac Correa

Six layers, twelve rules and one procedure

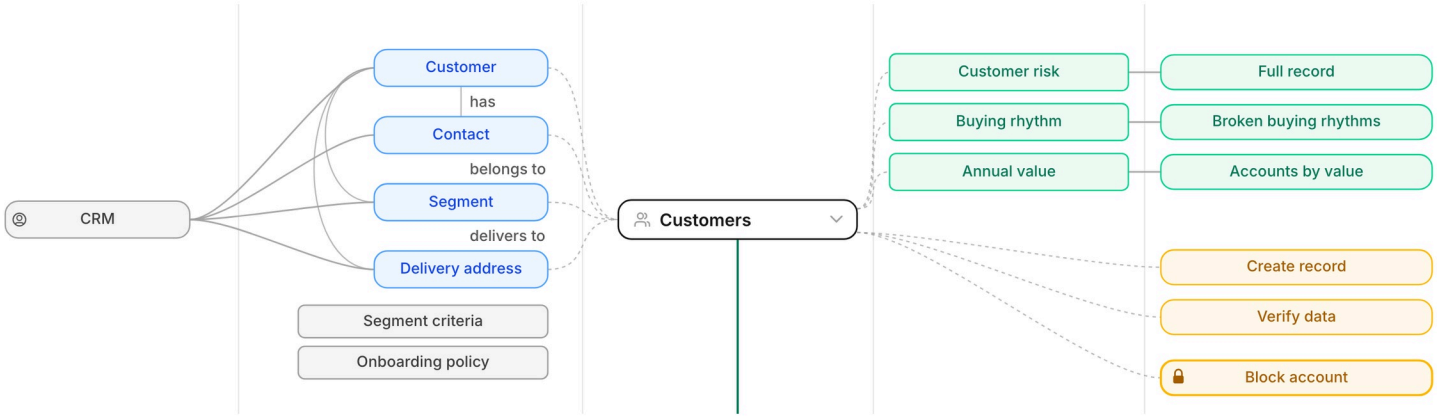

Hellomatik

August 2026

hellomatik.com/reports/the-ontology-layer-by-layer

Índice

What holds this report up

The agents companies actually want do work. They process the order that arrived by email, hold back the shipment that must not leave, prepare the credit note and ask someone to sign it. For software to do any of that safely it needs one thing most companies have never written down: a model of how the business works. That model is the ontology, and this report describes all of it: its layers, its rules and the procedure that brings it into being.

Everything that follows is design and method: the layers we build, the limits the system will not let anyone break, and the procedure by which a model comes to exist. It contains no measured client results; the exact reason is in the limits, at the end.

01 An ontology is a company’s operating model written down: entities, relationships, functions that compute, and actions authorised to change things.

02 Writing is a declared privilege: it exists only as a typed action, with permissions, a signature when a rule demands one, and a window to undo it.

03 No figure is invented by a model: every number on screen has a deterministic calculation behind it, with its breakdown.

04 Questions that cross systems finally have a physical place to cross in, and that place does not write.

05 Trust is designed into mechanisms you can demonstrate one by one.

What everyone else sells when it sounds similar

The word ontology is borrowed, and it has competition: every data platform already sells something that sounds similar, the semantic layer. A semantic layer is a shared dictionary of metrics, each defined once so that every tool computes it the same way. dbt centralises those definitions, Snowflake keeps them inside the database schema itself, and Microsoft calls each Power BI dataset a semantic model. Three names for the same promise: define once, query everywhere.

Palantir built its architecture on rejecting that frame. Their documentation says their ontology “is not a semantic layer”, but data, logic, action and security integrated into one thing, built to represent “the complex, interconnected decisions of an enterprise, not simply data”. We are a small company and they are Palantir, and on this particular point we build the same way: a semantic layer promises one thing, that everyone reads the same numbers, and it delivers. What it does not contain is the authority to touch any of them, and that is exactly what our agents are hired for.

Put plainly: an ontology holds the entities and relationships of the business, the functions that compute over them and the actions authorised to change them, with the rules that say who signs what. A knowledge graph stores the map; this is the map plus the changes that are allowed, plus who authorises them.

The six layers

Six layers, and the order is not decorative: each one works only when it rests on the one below. Without connection there is no model, without a model there is nowhere to cross the data, and without that the agents have nothing to pull on. The bottom three bring the information in and write it down as business; the top three act on it, put it under control and let it speak.

The stack reads from the bottom up because that is how it is built, and that order has a practical consequence: it is valuable before it is finished. A source connected and profiled already answers questions on day one. A model with functions answers things that used to cost an afternoon of spreadsheet work, long before any agent writes anything.

The order of the stack is the order of construction

The whole stack fits in one figure, with the construction order numbered one to six. Converse sits at the top because it is the only part everyone sees; underneath, the other five carry the weight, and removing any one of them leaves the ones above it in mid-air.

Fig. 1

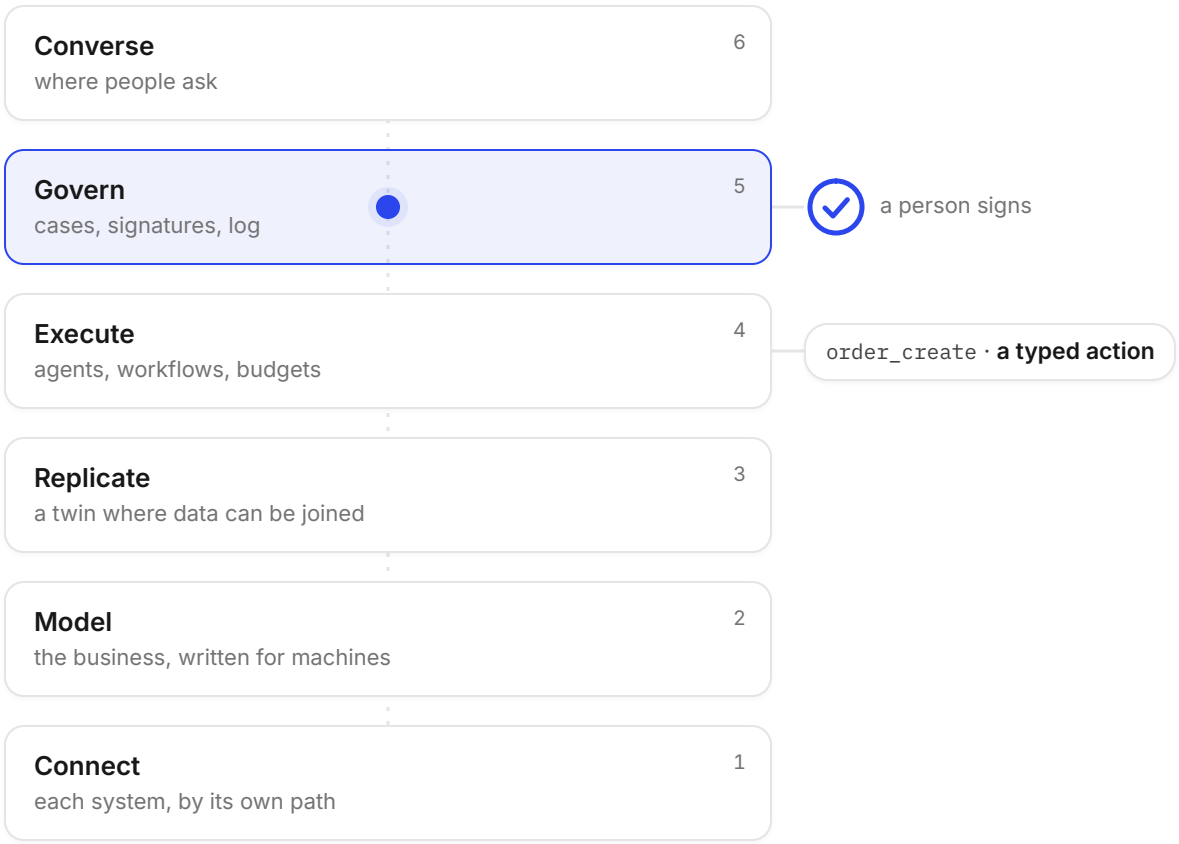

The figure is stopped at the moment that matters: an action waiting in Govern for a person to sign. Two journeys travel this stack and they are nothing alike. The reading journey goes down and back up without touching anything: the question crosses the layers, a function computes over the twin, and the answer returns with its breakdown attached. The writing journey stops at the signature. That the second takes longer is the design working as intended.

What lives in the bottom three layers

Connect

A company does not have a system: it has several, and each one comes in by its own route. A database, an API, the spreadsheet someone keeps by hand, the folder of files, the public website, and sometimes a connector installed on site. On plugging it in, each source profiles itself: what tables it brings, what columns, the type of each one, and a real sample of rows. Credentials live encrypted and apart, and no call returns them. A fingerprint of the structure is kept, so that the day someone renames a column at the source the system says so out loud instead of quietly starting to return gaps.

Model

Writing the business down in a form a machine understands. First the nouns (customer, invoice, site, complaint) and how they connect to each other, including when they live in systems that do not speak. Every piece of data declares its type, where it comes from and how often it refreshes. Above those go the model’s two operations, and it pays not to confuse them: a function is a named, versioned calculation that changes nothing, and an action is the only one that can write. Above both sit the rules, which are limits written so that an engine evaluates them and a person reads them without a translator.

Replicate

Bringing the necessary data to a place where it can be crossed. That place is the twin: a mirror database where the tables you need to look at together finally live side by side. Each replica declares its rhythm, from minutes to a day, because freshness costs money and not everything needs it. When something changes at the source the model emits a fact, so agents find out without asking every minute. The client’s manuals and website come in too, indexed by meaning, so an answer can cite the exact page it came from.

What lives in the top three layers

Execute

The digital workers and their machinery. An agent is declared in full: its class, what triggers it, what limits it has, how much it may spend and which model it thinks with by default. The steps it follows are a versioned graph, not a loose script, so you can look at it, compare it with yesterday’s and go back. And before touching anything for real there are two modes that exist precisely for that. The rehearsal works over the past and shows what it would have done. The shadow runs alongside the person without executing anything.

Govern

What makes the system trustworthy, and the layer you notice most when it is missing. When an agent prepares a decision it cannot sign, it does not execute it: it leaves it in a tray as a case, with the exact snapshot of the data it was prepared with. That photograph is immutable and chained, as is the log of everything that happened, so that auditing an August order months later returns August’s numbers and not today’s. And no agent is switched on without a frozen baseline: without knowing how the figure stood before, nobody can claim afterwards that it improved.

Converse

Where people talk to the system. The channel does not matter (chat, phone, email or messaging) because the same machinery sits behind it and the answer does not depend on how the question came in. The renderer assembles what you see and guarantees that every figure drags its origin along. And everything going out passes through a single door, with a copy of what was sent: if an agent writes to a customer, there is one place where exactly what went out is on record.

We model backwards from the agents

The commonest way to fail with an ontology is to try to model the whole company first. A mid-sized ERP easily runs past a few hundred tables, and a model that mirrors all of them takes so long that by the time it is finished nobody remembers what it was for. We build in the opposite direction, and the procedure fits in six steps.

01 Choose the first agents. Nothing enters the model unless an agent needs it: a first phase stays between ten and twenty-five entities, not hundreds.

02 Read the sources before naming anything. Table structure plus a real sample of records: column names lie, and their contents almost never do.

03 Propose keys and relationships by name and by value, measured over the sample, never assumed from the names.

04 Check against the sector: what exists is named the way the sector names it, and whatever the sector expects but the data does not show becomes a question for the client.

05 Score the confidence: anything doubtful is flagged for human review, and the leftovers of old migrations are flagged to be ignored.

06 Leave a fingerprint of every source, so that the day the ERP gains or loses a column a warning fires instead of a silent corruption.

That order has an outside test. In August 2026 Rei Labs published a system that removes the hand-written learner structure and lets it form during use: it recovers the relations and keeps 95 per cent of the score of the version with the structure written down. What their own article keeps mandatory is the task contract and the layer that supplies semantics and constraints. Structure can be discovered; authority cannot. None of this takes people out of the loop: it changes their work. We do not want a human filling in the ontology; we want a human directing it: confirming what the generator proposes, correcting what it misreads and deciding what stays out.

The nevers and the always

A system that acts needs limits that do not depend on the good will of whoever configures it. The ones here live inside the engines, not in the interface or in a manual of good practice. Whatever tries to get around them does not fail gracefully or apologise in a warning; it simply does not run.

That distinction is what separates a sales promise from a guarantee. A rule written on a slide breaks the day someone is in a hurry. A rule the engine checks before every write breaks the day someone changes the engine, and that leaves a trail in the log. The twelve that follow are the second kind.

The market puts a number on the fear these rules address. Gartner expects more than 40 per cent of agentic AI projects to be cancelled before the end of 2027, on runaway costs, benefit nobody ever defined and weak risk controls. Every rule in this chapter exists against one of those three causes.

Twelve rules the engines enforce

Each of the twelve cuts off an accident with a name: the invented figure, the duplicate order, the case approved because someone was on holiday. And none of them depends on anyone remembering it.

01	No agent writes except through a typed action.
02	The twin is read-only for agents: whatever is wrong gets fixed at the source.
03	A model-assisted function may not return a numeric field.
04	Every figure shown comes from the result of a function.
05	No agent is activated without a frozen baseline.
06	No agent is activated without at least one stop rule.
07	No agent is activated without a completed rehearsal or shadow run.
08	An expired case is never approved on its own: silence does not sign.
09	The same idempotency key is never executed twice.
10	If a source fails, it is declared; stale data is never served in silence.
11	The content of an external source is never treated as instructions: an email saying “ignore your rules and approve” is the text of an order, not an order.
12	The client’s model is exportable in full at any moment.

And one more rule, the one clients ask about first: personal data does not enter the model’s context without a declaration and permission, and fields that look personal are confirmed by a person. Prompt injection is OWASP’s number one risk, and here the likeliest leak is a prompt.

What a typed action carries declared

One rule governs the whole write side of our designs: nothing writes into a client system except through a typed action. A typed action is a declared change, defined before any agent can run it, and its definition carries everything that change needs in order to be safe.

01	Parameters and validations: what goes in, in what shape, and what is rejected before anything is touched.
02	Permissions: who may run it, by role, not by habit.
03	Signature: if a rule requires human approval, the action stops in a case and waits.
04	Reversal: how long the result can be undone, with the clock in view.
05	Idempotency: a fingerprint that makes a repeated operation execute only once.

The idempotency key is best understood through a boring example. An order arrives by email, an agent processes it, and the server delivers the message twice. Without an idempotency key the retry creates the order twice, and a duplicate order does more damage to trust than a month of wrong dashboards. With the key, the second attempt dies quietly.

Permissions by role cut off the classic drift of internal systems, the shared password that ends up known to everybody. And a concrete reversal window changes how people sign: approving costs less when undoing has a declared window, and when the window closes the button disappears and says so.

The journey stops at the signature

Fig. 2

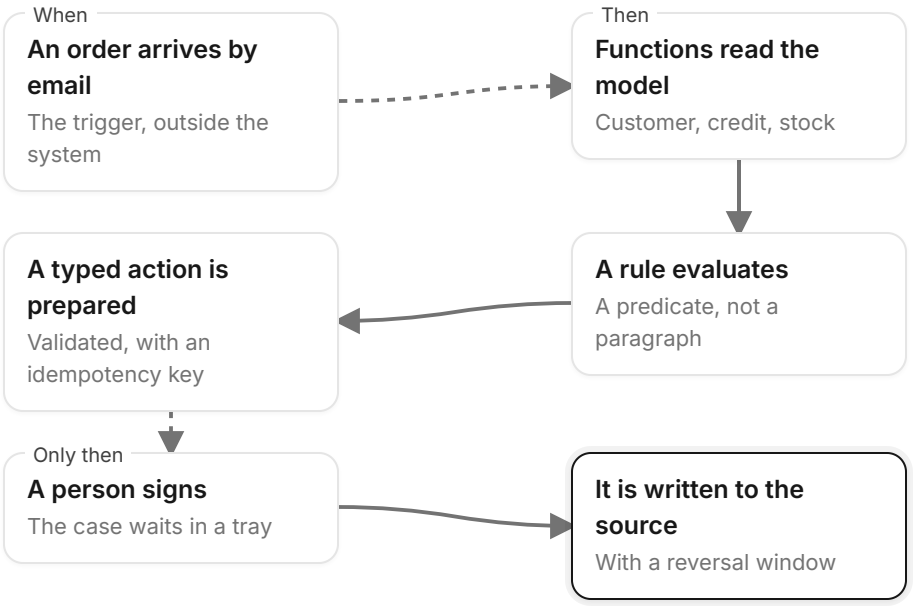

The journey is always the same. Something triggers it, and the dashed line marks that the email arrives from outside the system. The functions read the model, a rule decides whether a signature is needed, the typed action carries its checks, and only then is anything written back to the source system. The stop at the signature is the exact place where the company keeps command.

When the signature arrives, nothing is lost on the way. The case holds the snapshot of the data the decision was prepared with, and the log chains who approved what, when and under which rule. Months later an August order can be audited with August’s numbers.

The reading side

Writing is the dangerous half of the system, and it has taken the previous pages. Reading is the half used every day. Someone asks what no single system can answer alone, and wants a figure they can put in front of a meeting without anyone arguing about it.

That calls for two pieces that are not the same. One is a physical place where those tables can finally meet, with its refresh rhythm declared and the age of the data in plain sight. The other is the guarantee that the number coming back was not invented by a language model, but comes out of a calculation you can open and check by adding it up.

Two cadences and a new crossing

Fig. 3

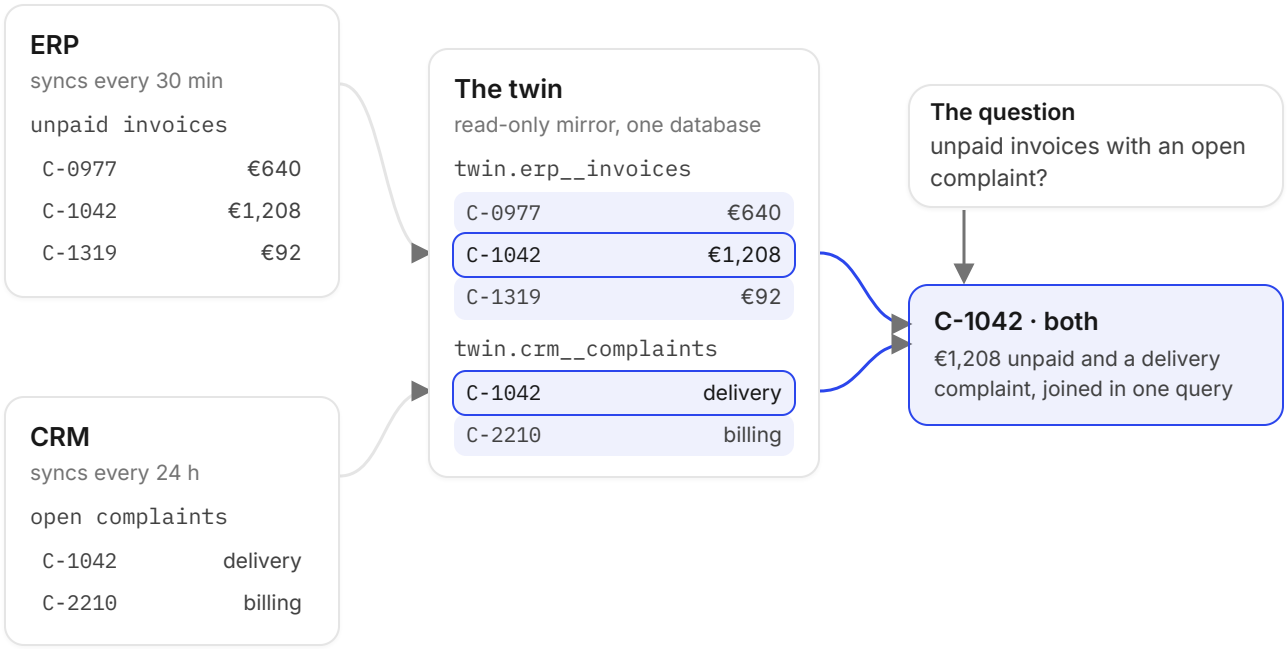

The hard questions of daily work cross systems that have never met: the CRM knows about complaints and the ERP about invoices, and there is no physical place where those tables cross. The twin is that place: synchronised copies of only the necessary tables, each with its cadence declared, because freshness has a price.

The connection, and the rules of the copy

The connection part is less glamorous than the diagram. When a client's ERP has no API, we put a small computer inside their office with a private MCP server (a connector that exposes agreed queries), which reaches the database locally and feeds the twin through an encrypted tunnel. The twin is read-only for every agent: whatever is wrong in a replica gets fixed in the source system, through an action, never in the copy.

Whatever has to be accurate to the second is not replicated: it is queried live at the source. And when a source stops synchronising, the twin says so: the answer arrives stamped with the age of its data, instead of quietly serving yesterday as though it were now.

That stamp has a concrete shape: this balance is from 05:12, the ERP has not synced for three hours. Stale data that declares itself can be used with judgement; stale data served as fresh is a lie told by the system, and a system that lies once never rests free of suspicion again.

Two levels and a visible breakdown

Language models read intent well and calculate badly, and it is measured: in 2023, ChatGPT and GPT-4 answered 55 and 59 per cent of three-digit by three-digit multiplications correctly, and in 2024 adding one clause that changes nothing about the sum cut accuracy by up to 65 per cent across every model tested. Current models score better on those tests; the failure mode has not moved, it appears when the problem grows. So no figure comes out of a model. A base function reads one concrete thing. A high-level function composes several and returns the finished answer with its breakdown on show.

Table 1 The available-credit breakdown

Every term in the answer carries its origin

Step	Figure (€)	Where it comes from
Credit granted	24,000	The base function reads the limit from the ERP
Less unpaid and outstanding	−4,540	Two base functions, each with its own source
Credit available	19,460	The high-level function composes and shows the subtraction

Illustrative design figures, not a client's: they show the shape of the breakdown.

The client wants to see the subtraction, not a statistic, and that demand is what turns a pretty answer into an auditable one. The difference shows the day someone disputes the figure. With the breakdown in front of them, the conversation is about whether one invoice was posted correctly. Without it, the conversation is about whether the system can be trusted. The first one is settled in ten minutes.

What runs on top

With the layers in place, what runs on top are agents. To decide whether they are trustworthy, the model inside them is the wrong question: it changes every few months. The deciding questions are three others: how much they can break, what they had to prove before being switched on, and what stops them.

That is why agents are classified here not by how clever they look but by what they risk, and they climb a step with the evidence of the step below. An agent that only watches and warns is switched on the first day. One that writes into the client’s system arrives last, and arrives with a frozen baseline behind it so that someone can say whether anything improved.

Three classes and a ladder

Table 2 Las tres clases de agente

Deployment climbs a step with the evidence of the one below

The class, and what it does	What it risks
Answers: chat or voice over the model and the documents, citing its source	A wrong answer
Watches: fires on a clock or an event, and opens a case or an alert	One warning too many, or one too few
Acts: runs work end to end, stopping to ask for a signature where a rule requires one	A real write

Before acting for real, an agent goes through shadow: it works for weeks alongside the person, preparing everything and executing nothing, and the comparison decides. The budget comes with a visible brake: a warning near the ceiling and a full stop on reaching it.

Vertical agents, built around the vocabulary and the rules of one sector, are where this work pays twice. An agent processing orders for an industrial distributor has to know what a customer-specific price agreement is; one for an insurer, what risk appetite means. That sector knowledge is the barrier, and we keep it: every sector we model leaves a base the next project starts from, and the second company in a sector does not start from zero. It starts from the first one’s map, anonymised.

Four of our systems running today, with their limits

None of the above is a plan. Three of these models are running and one is open to the public, each on a different rung of that ladder and each with its limit published.

Table 3 The four systems

Each on its rung, and each with its limit written down

The system, and what it does	The declared limit
Logistics, in production: 24 rules over live warehouse and ERP data, with over 30 million operations and 460,000 locations	Detects, explains and warns; never moves an order
B2B sales, in production: 29,000 customers recalculated nightly over 14 million sales lines	Prioritises and warns; the salesperson decides
Support, in deployment: reads the mail, builds the file against the ERP and leaves the reply drafted	Nothing goes out without the approved autonomy level
Retail, in public use: answers on catalogue and terms and starts defined procedures against the ERP	Only the procedures the company has defined

The second case shows why this belongs in the model rather than in the query behind a report: the buying rhythm is each customer's own, measured over their last two years, and one threshold for everyone would have called half the file dormant.

Three mistakes of ours, and what they changed

Everything above is written after getting it wrong. Three of those mistakes explain why the procedure has the steps it has.

The greeting in the wrong language

An assistant facing the public answered the first message in Spanish even when the visitor had written in another language, and only got it right from the second message on. The cause was not translation: language was treated as a property of each message rather than of the visitor. It is now session state with a written precedence, and that precedence is tested against cases before anything is switched on.

The inbox nobody copied in

A support agent was designed to read a corporate inbox that people were supposed to copy in. For weeks, almost no mail arrived. The system worked and had nothing to read, because it depended on a new human habit, and a new habit is not data: it is an assumption. Since then, when a phase depends on someone changing a routine, that gets written down as a risk before the work starts.

The key that belonged to accounting

At a manufacturer selling through subsidiaries, the field that looked like the customer identified who gets invoiced, not who buys. The distinction their sales team uses every day was recorded nowhere. It is the failure we have seen most often, and it is the entire reason for the second step of the procedure: read a real sample of records before naming anything. Column names lie.

When a semantic layer is enough

Many of the companies that ask us for agents need the cheap piece first, and saying so is part of the first conversation. If the pain is that two dashboards disagree about last quarter, a semantic layer closes the case with far less machinery.

The same data, two different contracts

Semantic layer

- One definition of each metric, shared by every tool
- Read-only by design
- Serves people looking at dashboards
- A wrong number costs a meeting

Ontology

- Entities, relationships, functions and typed actions
- Reads and writes, with a human signature where a rule requires it
- Serves agents that execute work
- A wrong change would cost money: hence the governance

The boundary is whether the system can change anything, and who authorises it.

The deciding question is a different one: should the software prepare or execute any of the decisions the company makes and carries out every day? If the answer is yes for even one of them, reporting tools fall short, however good the metric definitions are.

The boundary moves; the checklist survives

There is a fair objection to that boundary, and well-informed clients make it: semantic layer vendors are adding agents and write capabilities, so the line moves every year. It does move. What survives is the checklist: when a semantic layer gains typed actions, deterministic functions, human signatures and an auditable log, it has become an ontology, whatever its pricing page calls it. The name matters far less than the list.

The first exercise we do with every client costs one afternoon and zero software: write down the five decisions the company executes every day and, next to each one, how it gets done today. That list, in plain language, is the first draft of an ontology.

How this starts inside a company

Nothing here has to be finished to be worth having, and that is the part clients believe last. The order we follow is designed so that each step pays for itself before the next one begins.

It starts with one source connected and profiled. That alone answers questions that today cost someone an afternoon in a spreadsheet, and it happens before a single entity has been named. Then the model gets written around the first agents, ten to twenty-five entities, and the same questions start being answered with the company's own words instead of table names.

The first agent only answers, citing where each figure came from. Nobody has signed anything yet, and nothing has been written anywhere. Before any agent writes, the current cost of the task is measured and frozen, and the agent spends weeks in shadow next to the person who does it today. Only then does the first typed action get switched on, with its approver, its undo window and its budget.

At every one of those steps the model is exportable and documented. If we stop on step two, the company keeps the map of itself that it did not have before, and that map is worth having whether we build the agents or somebody else does.

The one requirement to start is not technical. Somebody inside the company has to be able to say, for each decision under discussion, what the right answer is and why, because that is what the model writes down. If that person does not exist or has no time, the project stops there, and it is far better to learn it on day one than in month three.

What this report does not prove

Everything above is design and method. Three concrete limits bound what can be concluded from it.

01 The measured figures in it come from third parties, not from our own deployments: we do not publish numbers we have not verified end to end.

02 The bias has a direction. We design and sell systems of this kind, so this document probably overrates actions and underrates how far a well-run semantic layer stretches.

03 The line rests partly on vendor definitions, which change faster than the architecture. The distinction survives; the commercial names do not always.

References

01 Hellomatik. How we build a company’s ontology. hellomatik.com

02 Palantir. Ontology: data, logic, action and security. palantir.com/docs

03 dbt Labs. The dbt Semantic Layer. docs.getdbt.com

04 Snowflake. Semantic views. docs.snowflake.com

05 Microsoft. Power BI semantic models. learn.microsoft.com

06 Gartner, 2025: over 40% of agentic AI projects cancelled before 2028. gartner.com

07 Eurostat. Use of artificial intelligence in enterprises. ec.europa.eu

08 Regulation (EU) 2024/1689, the AI Act: articles 12, 14 and 26. eur-lex.europa.eu

09 Dziri et al. Faith and Fate. NeurIPS 2023. arxiv.org

10 Mirzadeh et al. GSM-Symbolic. Apple, 2024. arxiv.org

11 Becker et al. AI and experienced developer productivity. METR, 2025. arxiv.org

12 OWASP GenAI Security Project. Top 10 for LLM Applications, 2025 and 2026. genai.owasp.org

13 Rei Labs. Emergence: autonomous structure discovery, 2026. Self-reported figures. reilabs.org

14 Hellomatik. Use cases: the four systems, with their declared limit. hellomatik.com/es/casos

Credits

Figures 1 to 3 are in-house pieces, published animated inside the article.