Mastercore Sinapsys Global® · May 2026

The Owner's Dashboard

Framework for SME management control in the era of agentic AI

AUTHORS

Freddy Arraez

CEO, Mastercore Sinapsys Global®

Aria by Mastercore Sinapsys Global®

AI Research Engineer

Buenos Aires · Argentina · tupymeclara.com

VERSION 1.0
May 2026

Executive abstract

Four questions every owner asks on Monday morning:

How much did I actually earn last month?
How much cash will I have eight weeks from now?
Is my budget on track, and if not, where is it slipping?
Which number do I have to move this week to save the year?

The **Latin American SME** operates within a mature ecosystem. Modern ERPs handle administration and statutory accounting with solvency. Tax and accounting professionals translate that compliance into legal certainty and fiscal strategy. Spreadsheets weave together whatever lies in between. Each actor plays its part.

On that solid base, however, one layer was never quite built: the **owner's dashboard**. A surface where the SME no longer asks *«am I compliant?»* but *«am I deciding well?»*. Until very recently, that layer was accessible only through expensive consultants or business intelligence teams that virtually no SME could afford to sustain.

Agentic AI changed that. Conversational language models with programmatic access to operational systems –via protocols like MCP–[1] now make something new possible: the owner *talks* to the system instead of operating it. But that conversation has value only if the system, on the other end, holds an explicit doctrine of what it means when it says EBITDA, what it means when it distinguishes cash flow from cash balance, what a favorable variance really is.

This whitepaper presents the framework of Tu PyME Clara® (Clarity SME): a system designed for that new layer. It has two sides. The **theoretical framework** makes the doctrinal decisions explicit –EBITDA without interest, accrual distinct from cash, variance read by impact on profitability– in language accessible to the owner and at the same time rigorous enough for the advisor auditing the result. The **pragmatic framework** materializes the doctrine in a multi-tenant, multi-country, multi-currency SaaS architecture, with native AI agent integration and doctrinal integrity validated through six rounds of external rigorous review.

We do not propose to replace the ERP or the accountant. We propose to complete a layer that the ecosystem needed and that technology has only just made possible.

Keywords: SME management control, managerial accounting, agentic AI, Model Context Protocol, EBITDA, accrual basis, cash flow, IFRS for SMEs, Latin America, Odoo.

[1] Model Context Protocol, proposed by Anthropic in 2024 as an open standard for communication between AI agents and external tools (Anthropic, 2024).

Contents

§1 · The ecosystem and the gap

> *After decades of excellent ERPs and accounting professionals, why does the SME owner still run the business with intuition and a spreadsheet?*

1.1 Three actors that deliver

The Latin American SME does not operate in a vacuum. Around it stands a mature, professional ecosystem that holds the business up.

The ERP. Odoo, SAP Business One, NetSuite, Microsoft Dynamics –and the dozens of regional solutions orbiting the same segment– solve, when well implemented, the administrative backbone: electronic invoicing, accounts payable and receivable, inventory, bank integration, reconciliation, statutory ledgers and regulatory reports. For a growing SME, installing an ERP is no longer a question of *whether* but of *which one*.

The tax and accounting professional. The figure without which no Latin American SME truly operates. They do far more than meet the tax calendar. They translate regulatory complexity into legal certainty, do fiscal planning, flag normative changes, sign the financial statements a bank will read to decide on a loan. In demanding jurisdictions like Argentina or Venezuela, that role is structural.

The spreadsheet. Excel, Google Sheets. Whatever the ERP does not contemplate, whatever the accountant does not need to see, whatever the owner wants to analyze in a particular way. The connective tissue that ties the rest together. An SME that claims not to use spreadsheets is hiding them better.

> A fourth actor is growing: BI tools (Power BI, Tableau, Metabase). Useful, but they require someone to maintain them. For an SME with fifteen or a hundred employees, that person rarely exists.

Each of the three actors does its job well. The easy critique –«ERPs are inflexible», «accountants only look backward», «spreadsheets don't scale»– ignores an uncomfortable truth: no one has a better version of what those actors solve today. Administrative operations, statutory compliance and ad-hoc flexibility have rightful owners.

1.2 The missing layer

What the ecosystem never had, however, is a layer above the three. A layer where the question is not *«am I compliant?»* but *«am I deciding well?»*.

The SME owner does not need to know, on Monday morning, what was the accounting entry for the last consulting invoice. They need to know whether their company made money this month, whether next fortnight they can cover payroll without tension, whether the budget variance is something serious –bank overdraft interest, for instance– or a one-off purchase already behind them.

That reading –managerial, not accounting; forward-looking, not backward– requires a different level of abstraction. It demands that revenue and expenses, accrual and cash, real and projected, be articulated into five or six indicators that the owner can interpret fluently. And it demands that the dashboard be available every day, not at the next meeting with the accountant.

Until very recently, that layer could be built in only two ways: by hiring an in-house controller –a professional dedicated to management control, whose salary a mid-sized SME cannot sustain permanently– or by paying an external consultant to build and maintain the dashboard. Both are valid, both exist, and both leave most of the segment outside.

1.3 Why now

The gap is historical, not structural. It is not that ERPs are hostile to management, or that accountants are wrong. It is that the management layer, until just months ago, required a technological infrastructure that has only just arrived.

Two advances enabled it, in parallel:

- **Language models** powerful enough to reason over structured data without collapsing into invention. Claude, GPT and the Codex peers already available in production.
- **Agent protocols** that allow those models to read and write on real systems with security and audit guarantees. MCP (Anthropic, 2024) is the first to consolidate as an open standard.

With both available, only in 2025 did it become technically possible to build a dashboard that the owner *converses* with instead of operating. A dashboard that understands questions like «*how do we close May at this sales velocity?*» and answers with the correct calculation, without hallucination and without inventing columns that do not exist.

Tu PyME Clara® (Clarity SME) is born at that intersection: in the layer that was missing, with the technology that has just arrived.

The ERP delivers. The accountant secures. The spreadsheet flexes.
What was missing was the dashboard that decides.

---

§2 · The thesis: theoretical and pragmatic frameworks

> *What does it take for a management system to serve the owner and, at the same time, to be taken seriously by the advisor who audits it?*

2.1 Two sides of the same product

Tu PyME Clara® (Clarity SME) is designed with two sides that must hold each other up.

The first side is the **theoretical framework**: an explicit accounting framework, accessible to the owner who never studied accounting and at the same time defensible to the management consultant who did. When Tu PyME Clara® (Clarity SME) says EBITDA, it means exactly *Earnings Before Interest, Taxes, Depreciation and Amortization* in the sense given by the *Corporate Finance Institute*[2]. When it says cash flow, it distinguishes movements from balances following the direct method of IAS 7 (International Accounting Standards Committee, 1992). When it says favorable variance, it reads the deviation by impact on profitability rather than by algebraic sign, aligned with the standard tradition of *variance analysis*[3].

That doctrine, explicit and citable, is what allows an external advisor to review the dashboard and sign at the bottom.

The second side is the **pragmatic framework**: a SaaS architecture that materializes the doctrine in a product used every day. Multi-tenant, multi-country, multi-currency, with a native agentic AI protocol. Designed so the owner can open the dashboard at nine on Monday morning, ask their AI advisor three questions and receive answers calculated on the same base a financial controller would have defended.

The two sides are not symmetric. Doctrine without product stays in a book. Product without doctrine fails the first audit. Tu PyME Clara® (Clarity SME) presents them together because only together do they solve the problem.

2.2 No accounting language, with accounting framework

There is a design choice worth declaring up front: **Tu PyME Clara® (Clarity SME) does not ask the owner to learn accounting.**

It does not use the word *debit* or *credit*. It does not talk about journal entries, chart of accounts, or analytical accounts. It does not require the user to understand double-entry bookkeeping.

But Tu PyME Clara® (Clarity SME) *is* accounting-rigorous on the inside. The separation between revenue and discounts respects IFRS 15 on revenue recognition[4]. Currency conversion in consolidation follows IAS 21[5]. The choice between direct and indirect method for cash flow follows the explicit recommendation of IAS 7 in favor of the direct method (International Accounting Standards Committee, 1992).

The short phrase is: *the owner's language on the outside, accounting framework on the inside.*

[2] CFI is the most-used reference in the Latin American finance industry for profitability metrics definitions. Its definition of EBITDA aligns with IFRS and US GAAP practice (Corporate Finance Institute, 2024).

[3] The favorable/unfavorable distinction is standard in university teaching of managerial accounting; an accessible treatment is available in the open educational resources of Lumen Learning (Lumen Learning, 2023) and in the professional reference AccountingTools (Bragg, 2024).

[4] IFRS 15 establishes that discounts and rebates are presented as a reduction of revenue, not as an expense. Tu PyME Clara® (Clarity SME) codifies this by separating the *income* and *discount* sections (IFRS Foundation, 2014).

[5] IAS 21 requires monetary items at historical cost to be translated at the closing rate of the reported period, not at the current rate (International Accounting Standards Committee, 1983). A common failure in management software is to use today's quotation and inadvertently re-express history.

Why this design

The hypothesis is that the owner does not need to learn the vocabulary of the balance sheet to make decisions about their business. They need to see six indicators on their dashboard and understand what each one means. Under that interface, the math can –and must– be as strict as any professional system. The simplification lives in the presentation, not in the calculation.

2.3 Agentic AI as a pillar, not an ornament

The third claim of the thesis is the most distinctive and deserves its own chapter further on (Section 6): **Tu PyME Clara® (Clarity SME) is designed with agentic AI as an architectural pillar, not as a feature bolted on afterwards**.

The difference is not rhetorical. A product with AI *added* exposes a side chat that answers what it can about the system's data, generally with poor results and frequent hallucinations. A product with AI *as a pillar* exposes an explicit contract –a surface of **skills**, in our case via the MCP protocol– that the external agent discovers, understands and operates under the same rules a human user would: permissions, validations, audit trail, idempotency, explanatory error messages. The AI agent, in this design, is not a client of the product: it is an *operator*, with the same rights and the same limits as a trusted employee.

That decision has concrete technical consequences detailed further on –structural anti-confabulation, idempotency by functional key, doctrinal validation on every write– and it becomes visible in the daily conversation between the owner and the AI advisor.

2.4 Complementary, not substitutes

It is worth closing this chapter with clarity about the position of Tu PyME Clara® (Clarity SME) against the ecosystem described in the previous chapter.

Tu PyME Clara® (Clarity SME) *does not replace* the ERP. It does not invoice, does not issue delivery notes, does not handle nominal receivables. If the client's ERP exports balances by concept and period, Tu PyME Clara® (Clarity SME) consumes them via AI agent, or by import, or by user upload.

Tu PyME Clara® (Clarity SME) *does not replace* the accountant. It does not sign financial statements, does not file taxes, does not advise on regulatory compliance. What it does is give the accountant –and the owner– a managerial reading of the same economic reality the accountant already knows, expressed in terms of decision.

Tu PyME Clara® (Clarity SME) *does not replace* the spreadsheet when an ad-hoc analysis is needed. What it offers is a structured alternative for everything that is *not* ad-hoc, the place where the spreadsheet used to stand alone.

That position –complementary, not a substitute– is deliberate. The SME ecosystem works because each actor covers their lane. Tu PyME Clara® (Clarity SME) adds a layer that did not exist and delivers it without asking anyone to give up ground.

Accounting framework inside, owner's language outside, agentic AI as the pillar. Three sides of the same thesis.

§3 · The five instruments of the dashboard

> *If I could look at only five things –and only five– to know how my business is doing this week, this month and this year, what would they be?*

3.1 The choice of five

A useful dashboard is not an exhaustive dashboard. It is a dashboard where each instrument answers a distinct question and no question is left without an instrument.

After nine years accompanying Latin American SMEs in their management processes –from Mastercore Sinapsys Global®, through ERP implementation, support and scaling on Odoo technology– we settled on five instruments that cover the full spectrum of managerial decision without redundancy. Each one answers a concrete question, operates on a distinct temporal dimension, and has an explicit accounting framework behind it (Table 1).

Table 1. *The five instruments of the Tu PyME Clara® (Clarity SME) management dashboard.*

Instrument	Question answered	Horizon
Income Statement	How much did the business actually earn last month, separating operating from financial?	Recent past
Balance Sheet	What does the business own and owe? What is my actual leverage?	Closing snapshot
Cash Flow	How much came in and out this month? Am I confusing movements with balances?	Recent past
Budget vs Actual	Is the plan being met? Do variances play for or against me?	Past vs plan
13-Week Forecast	Do I have cash to pay salaries and suppliers over the next three months?	Immediate future

We dedicate a paragraph to each below. The formal doctrine and design decisions behind each calculation are documented in the chapters that follow.

3.2 Income Statement: actual result, not statutory

The **managerial Income Statement** is the classic owner's question: *«did I make money last month?»*. It seems simple. It is not.

The most common trap is not in the calculation; it is in the *definition of the lines*. Specifically: what counts as operating expense and what counts as financial result. **EBITDA** –*Earnings Before Interest,*

Taxes, Depreciation and Amortization– exists precisely to separate operating profitability from the cost of financing (Corporate Finance Institute, 2024). If the interest on a bank loan gets posted as a «financial expense» inside operating expenses, EBITDA becomes contaminated. The owner believes they are seeing operating profitability and is actually seeing a hybrid that mixes how they operate with how they finance. Those are different things.

Tu PyME Clara® (Clarity SME) explicitly separates **bank fees** (operational, part of EBITDA) from **interest and financial costs** (non-operating, below EBIT). The user sees two categories; the system knows which each expense belongs to.

The Tu PyME Clara® (Clarity SME) Income Statement ends in five lines the owner can recite from memory: *net revenue, gross margin, EBITDA, EBIT, net income*. Each one with its percentage over sales. That is the read.

3.3 Balance Sheet: total leverage, not partial

The **balance sheet** answers a more structural question: what the business *owns* and *owes* at a given moment. For an SME, that reading matters for three things: to present to the bank if they are going to ask for a loan, to understand how much debt weighs on assets, and to detect accumulations that the month-to-month flow does not reveal (inventory that grows, receivables that stretch).

A common failure in SME management software is to present a *partial* balance sheet that omits bank cash from assets or live debt from liabilities, leaving the leverage ratio systematically underestimated. Tu PyME Clara® (Clarity SME) includes them because without cash there are no assets and without debt there are no liabilities. The bank's question –*how much do they owe over how much do they have?*– requires both sides complete.

3.4 Cash Flow: flow, not balance

The most important distinction in the dashboard is also the easiest to confuse: **cash flow is not cash balance**.

The balance is a snapshot: how much is in the bank at end of day. The flow is the film: how much came in and how much went out during the month. The two readings coexist and both matter, but they answer different questions. A system that sums them –that shows «total cash: $5,000,000» mixing the bank balance with the month's collections– delivers a number without economic meaning.

Tu PyME Clara® (Clarity SME) separates them structurally. The **cash flow** instrument of the dashboard shows the flow of the period: inflows, outflows, net. The bank balance lives in the balance sheet, where it belongs. The doctrine behind this is the **direct method** of IAS 7, which explicitly recommends this presentation for management reporting (International Accounting Standards Committee, 1992).

3.5 Budget vs Actual: the variance that matters

Budget vs Actual is where the owner discovers whether the plan is being met. The metric is the **variance**, and the trap is also in the reading.

A variance can be *algebraically positive or negative*. If we spent $100,000 more than budgeted on rent, the variance is +$100,000 over budget. Is that good or bad? It depends on the concept. On revenue,

spending more than budgeted is absurd (there is nothing to spend in revenue). On expenses, spending more is bad. Treating the variance algebraically –by sign alone– confuses the user.

The correct doctrine classifies variances by **impact on profitability**: **favorable** when they improve the result (more revenue or fewer expenses than the plan) and **unfavorable** when they worsen it (Bragg, 2024; Lumen Learning, 2023). That reading –the same any controller would use– lives encoded in Tu PyME Clara® (Clarity SME) at the level of the comparison view.

3.6 13-Week Forecast: the Monday decision

The last instrument –and the most operational– is the **thirteen-week cash projection**. Three months. The window where an SME makes its tension decisions: pay this supplier or that one, accelerate collection from a client or not, refinance the loan now or wait.

The thirteen-week forecast uses the **direct method** with weekly projection of inflows and outflows, anchored on the initial cash balance from the last real close. It does not replace the annual budget –which operates at a different level of abstraction– nor the three-year strategic plan. What it does is give the owner and the CFO –if there is one– the operational reading the rest of the dashboard does not give: when cash will tighten, when it will ease.

Tu PyME Clara® (Clarity SME) separates this instrument structurally because its calculation logic is different from the Income Statement: here it does not matter when the revenue is *accrued*, it matters when the *cash actually arrives*. The difference between accrual and cash –the heart of the next chapter– is what makes both instruments coexist without contradicting each other.

3.7 An implicit hierarchy

The five instruments are not equivalents. They have an implicit hierarchy worth making explicit: from the most operational –the weekly cash forecast– to the most strategic –the accumulated balance sheet– passing through monthly operations and the annual plan (Figure 1).

No instrument is redundant. An SME looking only at the base takes operational decisions without strategic context. One looking only at the peak loses reaction speed. The five are read together; each from its own horizon.

Five instruments, one hierarchy, one doctrine.
This is what the owner looks at to decide.

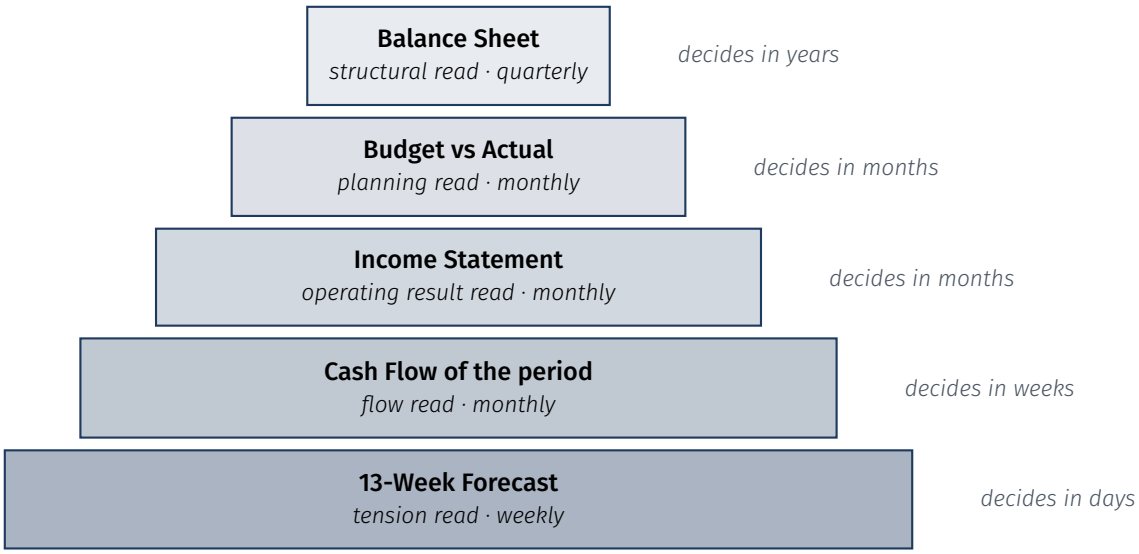

Figure 1. *Implicit hierarchy of the five instruments. The base –13-Week Forecast– is the most operational read, reviewed every Monday. The peak –Balance Sheet– is the structural read consulted at each quarterly close.*

§4 · The two dimensions of time

> *Why can the same expense appear in February, March or April depending on who looks at it, and why is that not a system error?*

4.1 The March consulting fee

Imagine an everyday situation. A consultant provides services to our SME during March. They issue an invoice on the 28th of the same month. The SME, orderly, pays at thirty days: April 15th.

That single economic event –a professional service received and paid– occupies three places in time from the owner's perspective. Service rendered in March. Invoice received in March. Payment effective in April.

For the owner who pulls the March Income Statement, the question is *«is this fee a March expense or an April one?»*. For the same owner looking at the next-fortnight Cash Flow forecast, the question is different: *«when did the cash go out?»*. Both questions are legitimate. Both have different answers. A good management system should be able to answer both without forcing the user to duplicate the information.

4.2 Accrual and cash as independent dimensions

Accounting framework resolves this tension with two concepts taught in the first year of any accounting program and that almost no SME software implements clearly: the **accrual** principle and the **cash** principle (also called cash basis).

The **accrual** principle recognizes revenue or expense in the period in which the economic event takes place –when the service is rendered, when the product is delivered, when the resource is consumed– regardless of when it is collected or paid. It is the principle that defines the result in the Income Statement, the balance sheet and any profitability read. It sits at the core of IFRS, US GAAP and nearly all modern accounting frameworks (IFRS Foundation, 2014).

The **cash principle** recognizes the movement only at the moment the money actually arrives or leaves. It is the principle that defines the Cash Forecast, the treasury position and the weekly read of financial tension.

The two dimensions must coexist. Without the first, the Income Statement lies. Without the second, the Cash Forecast lies. Mixing them into a single date field is one of the most common sources of error in SME management software, because it forces the user to choose between a correct read and another.

4.3 Three temporal dimensions in every movement

Tu PyME Clara® (Clarity SME) resolves this tension by encoding *three* temporal dimensions in every movement (Figure 2). Each one answers a distinct question and lives in a distinct field of the data model.

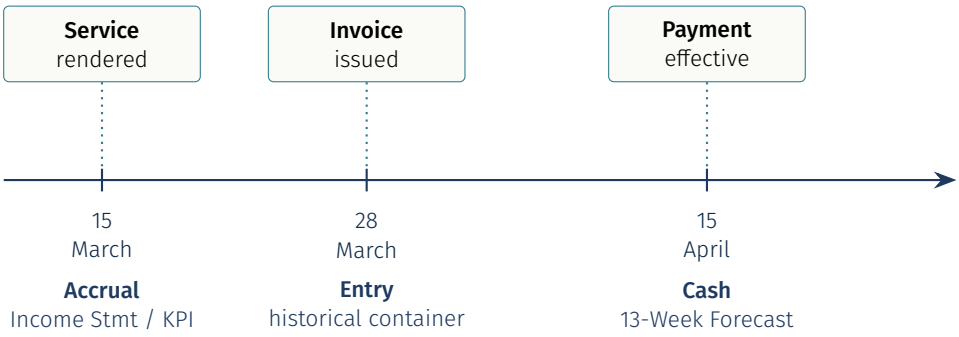

Figure 2. *Three temporal dimensions in every movement. The consultant's fee: service rendered on March 15, invoice issued on March 28, payment effective on April 15. The March Income Statement shows it (accrual). The Cash Forecast for the week of April 15 shows it too (cash).*

The short explanation of each dimension:

Accrual The period when the economic event is recognized. For the example fee: March. The Income Statement, profitability KPIs, Budget vs Actual comparison and balance sheet aggregate by this dimension.

Entry The period when the data was loaded into the system. For the example: it could be March (if the owner loaded it on receiving the invoice) or April (if they loaded it when paying). It is an operational container, not an economic dimension. It exists to maintain traceability of when the data entered.

Cash The exact date when money arrives or leaves. For the example: April 15. The 13-Week Cash Forecast and the period flow reading use this dimension.

The three dimensions can coincide –the simple case, where the owner pays at the moment of receiving the service– or differ. When they differ, each report still shows the truth from its own dimension, without contradictions.

4.4 What this resolves

This separation seems subtle but it resolves three classical pathologies of SME management software we see repeated in the market.

The monthly Income Statement that doesn't close. When a system uses a single date per movement, the owner ends up deciding whether the March invoice paid in April is «March expense» or «April expense». If they choose cash, the March Income Statement underestimates costs and overstates profitability. If they choose accrual, the April flow shows a ghost outflow. The usual solution in firms with well-kept books –playing with provision and accrual entries transactionally, then reversing them when cash actually moves– is accounting-valid but significantly raises administrative effort, and almost never reaches the owner's managerial dashboard.

The Cash Forecast that shows collections already past. Without a dedicated field for the cash impact date, the weekly projection inherits the disorder of the single-field design. Truly future collections appear mixed with collections already materialized. The grid loses operational value.

The comparison with the accountant that ends in argument. The accountant always works on accrual –it is the principle they sign in the statutory financial statements. If the managerial dashboard mixes criteria, the two readings are not comparable and the owner ends up explaining to the accountant why the system shows something different.

Tu PyME Clara® (Clarity SME) does not solve these pathologies by magic. It solves them by explicitly declaring that they are different problems that require different fields, and by putting the math so each report reads the correct dimension without forcing the user to choose.

On my dashboard, the month I sold is not the month I collected.
And both matter.

---

§5 · The sign doctrine

> *Why does a well-designed system never ask the owner to enter an expense as a negative number, and why is that not just usability but doctrine?*

5.1 The natural-entry rule

The owner who loads their office rent types a number. For instance: 450,000. They are not thinking about whether the number is positive or negative, and they don't care. They know it is an expense and they know its amount.

Tu PyME Clara® (Clarity SME) assumes the same logic. **The user enters magnitudes; the system assigns**

signs. If the concept is an expense –by operational definition– the system knows that 450,000 contributes negatively to the month's result. If the concept is revenue, it knows it contributes positively. If the concept is a management adjustment that can go in either direction, the system respects the sign loaded by the user.

This logic is encoded as an attribute of each concept that we call the **sign policy**. Three possible behaviors cover the operational reality of SMEs:

Always positive The concept always represents revenue. Examples: net sales, interest earned, recovered receivables. The system takes the absolute value of whatever the user loaded.

Always negative The concept always represents an expense. Examples: rent, salaries, taxes, bank fees. The system takes the absolute value and applies the minus sign.

Respects the loaded sign The concept can go in either direction depending on the case. Examples: management adjustments, foreign exchange differences, sale of decommissioned assets, financial results. The system respects the sign the user loaded.

The operational consequence is that the owner does not think about signs. They load reality as they live it: the rent is 450,000, not −450,000. When the report shows the result, the signs are applied correctly because the doctrine lives in the concepts, not in the data.

5.2 Variance: favorable, not algebraic

The sign doctrine has a second application that matters greatly to the owner and that almost no SME software solves well: how to read the **variance** in the Budget vs Actual comparison.

Let's go back to the rent example. Suppose the May budget was 450,000 and actual closed at 480,000. The algebraic difference is +30,000. The variance is +30,000. Is that good or bad?

If the concept were *net sales* –positive sign policy– spending 30,000 more than budgeted would be *additional revenue*: good. Since the concept is *rent* –negative sign policy– spending 30,000 more is *additional expense*: bad.

The same number, +30,000, has opposite readings depending on the concept. Treating the variance algebraically –by sign alone– asks the owner to remember what type of account each line belongs to and mentally invert the reading when needed. That is exactly what a good dashboard should avoid.

The correct doctrine classifies variances by **impact on profitability**. There are two possible readings, whatever the algebraic sign (Table 2):

Table 2. *Variance reading by impact on profitability. The* algebraic *column is what many systems show; the* economic *column is what the owner actually needs to read.*

Concept	Situation	Reading
Revenue	Actual > Budget	**Favorable**
Revenue	Actual < Budget	**Unfavorable**
Expense	Actual < Budget	**Favorable**
Expense	Actual > Budget	**Unfavorable**

The favorable/unfavorable distinction is not our invention. It is standard in university teaching of

managerial accounting and in the professional practice of management control (Bragg, 2024; Lumen Learning, 2023). What is our decision is to encode it at the model level –not at the screen level– and to apply it automatically without forcing the user to think.

5.3 Why it matters to the owner

There are two practical consequences of this doctrine the owner perceives directly.

Data entry is fast and free of sign errors. The user doesn't have to stop and think whether the amount goes with or without a minus sign. They enter the number and the system does the rest. An SME that enters two hundred movements per month saves hours of review and, more importantly, eliminates a classical source of operational error.

The dashboard reads without double interpretation. When the Budget vs Actual panel shows a line in green –favorable– the owner knows it played in their favor. When it shows a line in red –unfavorable– they know they lost against the plan. They don't have to look at the algebraic sign, don't have to remember what kind of account it is, don't have to invert anything. The reading is direct.

That sign doctrine and variance reading align is what allows a user with no accounting training to read a dashboard with the same clarity a professional controller would have.

The user enters magnitudes. The system assigns signs.
Variance is always read by impact on profitability.

---

§6 · Agentic AI as a pillar

> *What is the difference between a product «with AI» –like nearly everything today– and a product designed* from *agentic AI as an architectural pillar?*

6.1 Two categories that look alike but aren't

In 2024 and 2025, practically every management software product added an artificial intelligence layer. Side chats. Auto-summaries. Auto-completion. Market narrative says almost any product is *«AI-powered»*. The promise is heard everywhere.

But there are two very different categories under the same label and it is worth naming them.

Products with AI added. The product was designed five or ten years ago with a classical architecture –data model, UI, reports. AI was added later, generally as a side chat that looks at the system's data and tries to answer. Because it was not designed for that, the AI peeks at the database through the window:

it frequently invents columns, confuses periods, and when it makes mistakes the user has no way to know because there is no clear contract between the language model and the system. Hallucination, in this design, is structural.

Products with agentic AI as a pillar. The product was designed from the start so that an external agent –Claude, GPT, Codex and the peers that will come– becomes a legitimate *operator* of the system, not a client that observes. The agent has an explicit list of things it can do, validation of each action, an explanatory error protocol and permissions inherited from the user it represents. Hallucination, in this design, is blocked by construction: if the agent tries to do something that is not on its skill surface, the system rejects it before touching the database.

Tu PyME Clara® (Clarity SME) is in the second category. That decision is the most distinctive difference of the product, and this chapter argues why.

6.2 MCP: the governance protocol

The component that made the second category possible is relatively recent: the **Model Context Protocol** (MCP), proposed by Anthropic in 2024 as an open standard (Anthropic, 2024). The conceptual idea –as simple as it is powerful– is that a software system publishes a *catalog of skills* that an external agent can discover, understand and execute.

Each skill has four explicit things:

- A **name** and a description in natural language (the agent understands what it is for).
- An **input schema** (what parameters it accepts, with types, examples and valid values).
- An **output schema** (what the system returns).
- An **error contract** (what happens when something is wrong, with explanatory message).

The external agent does not write SQL nor manipulate the database directly. It asks the system, via MCP, to execute a determined skill with determined parameters. The system validates, executes, returns result or error. If the agent tries something off the list, there is no way for it to succeed: the protocol rejects it before any write.

Tu PyME Clara® (Clarity SME) currently exposes around thirty MCP skills, grouped in five categories: workspace state, management data, executive diagnostic, cash forecast, and plan/budget. Each skill has its embedded doctrinal validation –a movement that violates the EBITDA/interest separation, or the flow/balance distinction, or the functional-key idempotency, is rejected at intake, not discovered in the report.

6.3 Anti-confabulation as a design decision

There is a technical term from the AI agents literature worth naming: **confabulation**. It is the natural tendency of a language model to generate plausible answers that are not supported by real information. In trivial conversational contexts it has no consequences. In SME management systems –where a confabulated answer can lead to a wrong financing decision– it is unacceptable.

Tu PyME Clara® (Clarity SME) incorporates an explicit set of anti-confabulation mechanisms that live in the protocol, not in the model:

- **Closed skills.** The agent can only do what the system declares it can do. There is no back door via SQL or via generic API.
- **Doctrinal validation on every write.** When the agent tries to load a movement, the system validates that it respects the rules already documented: coherent sign, valid period, flow/balance separation, etc. If something doesn't fit, the system rejects with a message the agent can interpret and correct.
- **Idempotency by functional key.** If the agent tries to load the same movement twice –a classical conversational failure where the model gets confused and repeats– the system detects it and replaces rather than duplicates.
- **Verbatim summaries to the owner.** Every action the agent executes is narrated to the user in natural language with the real numbers the system wrote. If the narration does not match the data, the owner notices.

This combination –closed protocol, doctrinal validation and idempotency– blocks the three error families seen in classical chatbots tied to legacy systems: field hallucination, duplicated operations, inflated narrative that does not match what actually happened in the database.

6.4 A real conversation

To make the difference concrete, it is worth describing how a typical conversation looks between the owner and their AI advisor in Tu PyME Clara® (Clarity SME).

Monday morning. The owner opens the product's chat and writes: *«how did we close April? I want an executive analysis and, if there is cash tension in May, a recommendation.»*

The advisor answers with a close diagnostic: net revenue, EBITDA, EBIT, monthly result, comparison with March and with budget. What it says is not invented: every number comes from a dashboard skill – `EXECUTIVE PANEL READ`, `COMPARISON AGAINST BUDGET`, `COMPARISON AGAINST PREVIOUS PERIOD` – that the system executed with doctrinal validation. The narrative summary is from the language model; the data is from the system.

Next, the advisor consults the thirteen-week forecast via the skill `CASH PROJECTION`. It detects a tension in the second half of May –a tax outflow coinciding with a month of slower collections. It recommends two actions: accelerate invoicing of three clients with thirty-day terms to shorten them to fifteen, or arrange a bridge credit line with a specific financial provider.

The owner leans toward the first. They ask the advisor: *«load this as an adjusted budget scenario for May, so I can see how the forecast looks if I actually manage to shorten the collection cycle».* The advisor invokes the skill `LOAD ON BUDGET SCENARIO`, records the three adjustments with attribution to the user who authorized them, and the thirteen-week grid recalculates instantly. The second half of May now appears relieved. The owner decides to make the calls with that projection in hand.

Fifteen minutes of conversation. Zero clicks on forms. The decision that matters, made with the correct information.

6.5 Why this design was only possible now

It is worth closing the chapter with chronological clarity. What has been described was not technically possible three years ago.

The language models powerful enough to reason over thirty skills with structured schemas, without confusing a parameter, without inventing results, reached the market between 2024 and 2025. The MCP protocol, which lets those models talk to external systems with security guarantees, became standard in the same period. Earlier, attempting this design would have produced a frustrating product: the agent hallucinated, the system couldn't protect itself, the user ended up returning to the classical form.

Tu PyME Clara® (Clarity SME) did not anticipate this technology. It waited for it to mature and incorporated it as an architectural pillar. That is the difference between *product with AI added* and *product designed for the agentic era*. The first category has existed for years; the second is just beginning.

The agent is not a client of the system. It is an operator with permissions, validations and traceability. Like any trusted employee.

---

§7 · Doctrinal integrity

> *How does the owner –and the advisor who will review the dashboard– know that the doctrine the product declares is actually implemented underneath?*

7.1 Stating the doctrine is not the same as living it

A product can claim on its website to separate EBITDA from financial costs, to distinguish flow from balance, to apply IAS 21 on consolidation. That a product *says* so is one thing. That it *delivers* in every calculation, in every report, in every response from the AI agent, is another. The distance between stated doctrine and implemented doctrine is where silent problems are born: the EBITDA contaminated by interest, the dashboard that mixes balances with collections, the variance read backward.

To close that distance, Tu PyME Clara® (Clarity SME) submitted –and continues to submit– to a process of external rigorous review that documents the doctrinal integrity of the system line by line. We describe it in this chapter not because it is a technical detail, but because it is the operational difference between a product *that states* and a product *that demonstrates*.

7.2 Rigorous review by external AI agent

The method is simple and worth declaring. Tu PyME Clara® (Clarity SME) delivers its full source code –models, reports, MCP skills, data migrations– to an independent external AI agent, with no access to previous design decisions, configured as a senior auditor of Odoo Engineering and managerial accounting. The agent reviews the code and produces a forensic report where every finding comes with

file, line number, evidence, concrete risk and suggested fix.

The auditor we use is Codex, the OpenAI agent specialized in code review. The choice is not casual: using an agent external to the family of models that already operate the system (Claude, in our case) reduces self-validation bias. Codex audits without knowing the decisions taken previously; it judges the code as a rigorous human reviewer would find it, without diplomacy.

7.3 Six passes, thirty-five findings

The process was not linear. The first pass revealed twenty findings –four critical, eight high, four medium, four low– with an explicit verdict of *«not suitable for external review»*. Each finding pointed to a concrete distance between what the system said and what the code did.

After five sprints of correction and five additional re-audits –each one adding new findings as the depth of analysis increased– the final balance was **thirty-five findings audited, thirty-three closed, two deferred by explicit architectural decision** (Table 3).

Table 3. *Synthesis of the six Codex rigorous review passes. Deferred decisions are: the deliberate rejection of an annual immutable close –the SME often needs to adjust months after closing– and a doctrinal glossary that this whitepaper now covers.*

Pass	C	H	M+L	Status
1. Original audit	4	8	8	Not suitable: structural gaps in EBITDA, scenarios, cash.
2. Re-audit v2	4	4	0	Not suitable: incomplete coverage on secondary surfaces.
3. Hotfix v2	1	2	0	Not suitable: idempotency misaligned in Cash Forecast.
4. Hotfix v2 final	0	1	1	Not suitable: wizard aggregation without sign normalization.
5. Minor Sprint audit	0	2	0	Not suitable: accrual incomplete in MCP dashboard and public share.
6. Minor Sprint hotfix	0	0	0	**Suitable for external review, final.**

The two deferred decisions deserve mention. The first –not implementing an annual immutable close– is deliberately contrary to formal accounting practice, where the year's balance is signed and not touched. The SME, by contrast, adjusts findings months after closing –a late invoice, an imputation error discovered, a refund that appears– and forcing immutability would invite a parallel spreadsheet system. Our doctrine is **perpetual soft close**: the close remains as a reference, corrections are recorded with explicit audit trail, the report visibly flags when a line was modified after the original close. It is an SME decision, not an omission.

The second –the doctrinal glossary– is precisely what this whitepaper and its Appendix A now cover.

7.4 What this means for the owner and for the advisor

The honest conversation has two sides.

For the owner. It means every number on the dashboard has demonstrable backing. When the Tu PyME Clara® (Clarity SME) panel says *«EBITDA for the month: $3,250,000»*, that is not a figure obtained with an approximate calculation. It is exactly the difference between net revenue and operating expenses for the accrued period, without interest, without amortization, validated against the CFI/IFRS definition and audited by an independent reviewer. The promise of the banking-grade dashboard is

not rhetoric: it is a demonstrable property of the App.

For the advisor reviewing the dashboard. It means the reports are defensible. An accountant, a management consultant or a controller stepping in to evaluate their client's information can ask for the detail of the doctrine –how each line is defined, what is included and excluded, what design decisions were made– and receive answers verifiable against the App's documentation and against the accounting frameworks cited. There is no *«it's that way because we made it that way»*. There is *«it's that way because IAS 7 recommends it, IFRS 15 supports it, and the implementation was audited in pass 4 of the Codex forensic process»*.

That difference is what separates a system the advisor tolerates from a system the advisor recommends.

7.5 The method continues

It is worth closing this chapter with clarity about the future. Rigorous reviewing is not a milestone achieved once and kept as a badge. It is a continuous process. Every new sprint of the product –a new MCP skill, a change in the KPI engine, a data migration– passes through at least one external audit pass before going to production. The commitment is explicit: no change touches client management code without documented forensic review.

This practice has a cost –in iteration time and release agility– and we take it on as structural investment. An SME management product that is not audited loses, over time, the property that matters most: the trust of the owner and the advisor.

If the Agent claims it, the system can defend it.
Concept by concept, calculation by calculation.

---

§8 · SaaS architecture for Latin American SMEs

> *The textbook SME does not exist. The real Latin American SME does: with two countries, three functional currencies, an exchange rate that moves every week, formal invoicing coexisting with unbilled operations that also need to be measured, budgets that inflation makes obsolete by mid-year, and a fiscal close that gets adjusted months later because that is how the context operates. How does a management system fit that reality?*

8.1 What the theoretical framework needs from the pragmatic one

The doctrine made explicit in earlier chapters –the EBITDA/interest separation, the accrual distinct from cash, the variance read by impact on profitability– is portable. It applies to a Chilean professional services SME and to an Argentine group with three business units across provinces. It applies to a

commerce that operates only in local currency and to a company that invoices in dollars and pays salaries in local currency.

For that doctrine to be actually operable, the product architecture must respect SME reality as it exists, not as it appears in textbooks. This section describes the pragmatic framework features that make this adaptation possible without forcing the owner to deform their business to fit the system.

8.2 Multi-business with native consolidation

A significant proportion of Latin American SMEs operate through *more than one legal entity*. The owner has one company that invoices, another that owns the building, a sole proprietorship for some professional services. Economically it is one business; legally it is three.

Tu PyME Clara® (Clarity SME) treats each unit as an independent **business** with its own **chart of concepts** –the SME's *alphabet*–, its own functional currency and its own period calendar. At the same time, it allows defining **consolidations** –explicit combinations of units, with participation percentages and conversion rules– that produce a unified economic reading. The owner can look at «*my operating company*» with one click and at «*the whole consolidated group*» the next, without reloading anything.

8.3 Multi-country, catalog by jurisdiction

The standard chart of accounts for an Argentine SME is nothing like that of a Venezuelan one. The first has Gross Income Tax, Health and Safety Tax, Bank Debit/Credit Tax, year-end salary bonus, Employer Contributions, work-risk insurance, Income Tax, Personal Assets Tax. The second has IGTF, Municipal Tax on Economic Activities, Social Security, LPH housing tax, INCES, Income Tax. There is very little overlap.

Tu PyME Clara® (Clarity SME) solves this with a **declarative catalog**[6]. Each concept in the catalog declares a list of countries where it applies: empty if universal, `[AR]` if Argentine-only, `[VE]` if Venezuelan-only. When a workspace is configured for a specific country, it sees only the concepts relevant to its jurisdiction.

For a multinational company with units in different countries, each unit has its localized catalog, and consolidation works in the management currency of the group.

8.4 Multi-currency with closing-date exchange rate

Three different currencies can coexist in a single Latin American SME. The **functional currency** of each unit (Argentine pesos, Venezuelan bolivars, US dollars). The **source data currency** (sometimes the operator loads an invoice in USD even though the company operates in pesos). The **consolidation currency** of the group (at holding level).

Tu PyME Clara® (Clarity SME) codifies the three dimensions. Each movement is loaded in its source currency with its explicit exchange rate; it is persisted in the unit's functional currency; it is converted to the consolidation currency when the report requires it. The conversion follows the IAS 21 recommendation (International Accounting Standards Committee, 1983): exchange rate at the close of the reported period, not at today's rate. This avoids a common error in SME management software –re-expressing

[6] The curated concept catalog is published as a versioned configuration file, outside the calculation code. This declarative architecture lets local fiscal advisors propose catalog extensions without requiring a development cycle.

history with the current quote– that distorts temporal readings in jurisdictions with high currency volatility.

8.5 Hierarchical segments

The textbook SME has a single line of business. The real SME usually has several: geographic branches, sales channels, product families, service lines. And often those dimensions cross: the owner wants to see «*revenue by branch*» but also «*profitability by service line within each branch*».

Tu PyME Clara® (Clarity SME) solves this with a **hierarchical segment** system configurable per unit. The owner defines the structure that reflects their business –without imposing a predefined model– and reports aggregate by root, drill down by descendant, or pivot as needed. The hierarchy respects consolidation at the right level: when a parent segment has loaded children, reports sum without duplicating.

8.6 Perpetual soft close

Formal accounting practice closes the fiscal year and signs. After closing, entries are not touched. Any subsequent correction goes to a later period.

The Latin American SME does not operate that way. A late invoice appears in May for a service rendered in March. A supplier sends a credit note three months later. A customer disputes a sale and it gets adjusted retroactively. The owner *needs* to correct past facts when they appear, not at the next meeting with the accountant.

Tu PyME Clara® (Clarity SME) adopted an explicit architectural decision on this: **perpetual soft close**. The period gets marked as *closed* –which enables its export to executive reports, comparative KPIs and consolidations– but does not become immutable. If correction is needed, correction happens, and the system leaves visible audit trail of what was modified, when, and why. Readings that depend on the close recalculate automatically.

This decision is deliberately contrary to formal accounting practice of the immutable annual close. We adopt it because the alternative –forcing immutability– pushes the owner to keep a parallel system in spreadsheets, exactly what the dashboard comes to avoid.

8.7 Traceability as a structural property

The five features above share a property worth declaring explicitly: **traceability**.

Every movement in Tu PyME Clara® (Clarity SME) has **attribution** (who loaded it, when, and who was the last to modify it), **explicit origin** (from what source it came: portal, import, AI agent, manual adjustment), **access logging** (each read, write or deletion is recorded in an internal log with user and date), and **doctrinal context** (to which concept, in which accrual period, with what cash impact date, in which segment, in which scenario). When an external advisor receives the client's dashboard, they can ask for the detail of any line and receive a verifiable answer.

This traceability is not a feature added for compliance; it is a structural property of the data model. Without it, the doctrinal integrity of the previous chapter could not be sustained. With it, the system stays ready for external review without additional effort.

The SME does not fit into the system. The system fits into the SME and the jurisdiction where it operates.

§9 · Closing: doctrine, agentic AI, architecture

> *What changes, in the owner's life, when the dashboard is built with explicit doctrine, AI agent as a pillar, and architecture designed for Latin America?*

9.1 The TPC method in three lines

If we had to summarize all the content of the previous chapters in three lines, these would be them.

One: **the owner's language on the outside, the accounting framework on the inside**. The user loads reality as they live it –magnitudes, dates, everyday business concepts– and the system applies EBITDA per CFI, accrual per IFRS, FX per IAS 21, *variance* per the standard teaching of managerial accounting. Rigor lives in the code, not in the interface.

Two: **agentic AI as an architectural pillar, not an ornament**. The product was designed from the start so an external agent –Claude, GPT, Codex– operates the system like a trusted employee, with permissions inherited from the user, doctrinal validation of every action, and explanatory error messages. Hallucination is blocked by construction, not controlled by retroactive patches.

Three: **an architecture that fits the Latin American SME as it exists**. Multi-business with native consolidation, multi-country with catalog by jurisdiction, multi-currency with closing-date exchange rate, hierarchical segments that respect operational reality, and perpetual soft close because the SME adjusts months after closing and the system should not impose a rigidity that formal accounting demands but daily operation denies.

The three lines are not alternatives. They operate together or they do not operate. A doctrine without native agentic AI stays in a textbook. An AI agent without explicit doctrine confabulates elegantly. A Latin American architecture without the first two misses the historical opportunity.

9.2 What we say to the ecosystem

It is worth closing by clearly acknowledging the position of Tu PyME Clara® (Clarity SME) before the actors already covering SME operations –the ERPs, the accounting professionals, the spreadsheets, the BI solutions.

To the ERPs –with which at Mastercore Sinapsys Global® we work through implementation, support and scaling on Odoo technology, and from which we learned what we know about SME back-office– we say their operational role is structural and Tu PyME Clara® (Clarity SME) does not come to dispute it. When an ERP exports closing consolidated balances, Tu PyME Clara® (Clarity SME) consumes them. When not, the data is loaded via AI agent, import, or user upload. The boundary is clear: the ERP is the system of

administrative execution; Tu PyME Clara® (Clarity SME) is the managerial dashboard that draws from that execution to produce decision-grade readings.

To the accounting professionals we say their fiscal-regulatory advisory role is irreplaceable and Tu PyME Clara® (Clarity SME) does not sign financial statements, does not file taxes, does not advise on compliance. What it does is give the advisor a defensible managerial dashboard, audited by external reviewer, with explicit and citable doctrine. For the accountant stepping in to evaluate their client's information, Tu PyME Clara® (Clarity SME) should feel like a natural extension of their judgment, not a competitor.

We do not dismiss spreadsheets. Excel and Google Sheets are –and will remain– the most flexible ad-hoc analysis tool in existence. Tu PyME Clara® (Clarity SME) offers a structured alternative for everything that is *not* ad-hoc; pure ad-hoc stays in the spreadsheet where it belongs.

To traditional BI solutions –Power BI, Tableau, Metabase– we acknowledge their analytical power and their place in companies with a dedicated BI team. Tu PyME Clara® (Clarity SME) targets the SME segment where that team does not exist; when a company grows to a size where BI makes sense, Tu PyME Clara® (Clarity SME) can coexist with them without friction, exporting its stable metrics as a reliable source for analytical dashboards.

9.3 What we offer the owner

For the SME owner who reached this point, the concrete offering is simple.

A dashboard read in five minutes on Monday morning, with the five instruments covering the full spectrum of managerial decision: the result of the past month, the accumulated balance sheet, the period's cash flow, the budget comparison, and the thirteen-week projection that defines operational tension.

An AI advisor to converse with in natural language about the dashboard's own information, capable of diagnosing the close, proposing plan adjustments, simulating scenarios and loading new movements with doctrinal validation. Not a chat that guesses; an agent that operates the system under the same rules as a trusted employee.

An architecture that respects the real SME –multi-business, multi-country, multi-currency, with hierarchical segments and perpetual soft close– without asking the owner to deform their business to fit the system.

Doctrinal integrity validated by six rounds of external rigorous review, where the system can defend, concept by concept and calculation by calculation, where every number comes from.

9.4 And at the end, Monday morning

When we started this whitepaper we wrote the four questions the owner asks on Mondays:

How much did I actually earn last month? How much cash will I have eight weeks from now? Is my budget on track, and if not, where is it slipping? Which number do I have to move this week to save the year?

All four are legitimate and all four deserve answers calculated with doctrine, not improvised by intuition. Building a dashboard that answers them at this level of rigor was, until very recently, an expensive enterprise reserved for corporations with dedicated management control teams. The combination of

explicit doctrine, native agentic AI and Latin American architecture makes it accessible to the segment that needs it most.

That is the bet of Tu PyME Clara® (Clarity SME), and that is the editorial content Mastercore Sinapsys Global® brings to the table of owners and to the community of advisors and consultants who accompany them.

The dashboard exists. The doctrine is written.
The conversation with the system starts next Monday.

---

§A · Doctrinal glossary

Canonical definitions of the doctrinal terms used in the body of this whitepaper. Each entry includes the authority supporting it. The glossary is intended so an external advisor –accountant, management consultant or controller– can validate the terminology the dashboard uses against accounting frameworks and professional literature.

EBITDA — *Earnings Before Interest, Taxes, Depreciation and Amortization*. Operating result before financial costs, income taxes, depreciation and amortization. The most-used metric for pure operating profitability. (Corporate Finance Institute, 2024)

EBIT — *Earnings Before Interest and Taxes*. Operating result after amortization and before financial costs and taxes. Equivalent to *Operating Income* in US GAAP terminology.

EBT — *Earnings Before Taxes*. Result before income tax. EBT = EBIT − financial result ± non-operating results.

Net income — Result after income tax. The last line of the statutory Income Statement.

Accrual — Principle that recognizes revenue and expenses in the period in which the economic event takes place, regardless of when payment or collection occurs. It is the governing principle for the Income Statement and Balance Sheet in IFRS and US GAAP. (IFRS Foundation, 2014)

Cash (basis) — Principle that recognizes the movement only when money actually arrives or leaves. It is the governing principle for the Cash Forecast and operational treasury. IAS 7 doctrine recommends the *direct method* for management reporting. (International Accounting Standards Committee, 1992)

Favorable variance — Deviation of actual from budgeted value that improves the result: more revenue or fewer expenses than planned. Read by *impact on profitability*, not by algebraic sign. (Bragg, 2024; Lumen Learning, 2023)

Unfavorable variance — Deviation that worsens the result: less revenue or more expenses than planned.

Same impact-reading doctrine.

Sign policy Attribute of each concept that determines how the amount loaded by the user is signed: *always positive* (revenue), *always negative* (expenses), or *respects the loaded sign* (adjustments and ambivalent results). Lets the user load magnitudes while the system assigns the correct economic sign.

Perpetual soft close Architectural decision of Tu PyME Clara® (Clarity SME): a period is marked *closed* to enable executive and comparative reports, but does not become immutable. Later corrections –a late invoice, an adjustment– are recorded with visible audit trail. Differs from the *hard close* of formal accounting.

Accrual, entry, cash The three temporal dimensions encoded in every Tu PyME Clara® (Clarity SME) movement: the period of economic recognition (*accrual*, feeds Income Statement and KPI), the period when the data was loaded (*entry*, historical container), and the exact date when cash arrives or leaves (*cash*, feeds Forecast).

Flow *vs.* stock Doctrinal distinction between *cash movements* (period inflows and outflows, feed 13-Week Forecast) and *cash balances* (snapshot at close, feed Balance Sheet). Mixing them in the same report produces numbers without economic meaning.

Agentic AI Category of artificial intelligence where the language model not only answers text but *operates* external systems via skill protocols. Unlike the assistive chat, the AI agent has permissions, audit and doctrinal validation on every action it executes on the system.

MCP *Model Context Protocol*. Open standard proposed by Anthropic in 2024 for software systems to publish a catalog of skills that external AI agents can discover, understand and operate with security and traceability guarantees. (Anthropic, 2024)

Skill Unit of action a system publishes via agentic protocol. Each skill has a name, natural-language description, input schema, output schema, and error contract. The AI agent can only execute what is declared in the skill catalog.

Anti-confabulation Set of architectural mechanisms that block by construction the tendency of language models to generate plausible answers without backing. In Tu PyME Clara® (Clarity SME): closed skills, doctrinal validation on every write, idempotency by functional key, and verbatim narration to the user.

Traceability Structural property of the data model: every movement has attribution (who, when), explicit origin (source: portal, import, AI agent, manual adjustment), access log (reads, writes and deletions logged) and complete doctrinal context (concept, accrual period, cash date, segment, scenario).

§B · References

Anthropic (2024). *Introducing the Model Context Protocol*. https://www.anthropic.com/news/model-context-protocol. Especificación abierta para comunicación entre agentes IA y herramientas externas.

Bragg, S. (2024). *Variance analysis definition*. https://www.accountingtools.com/articles/variance-analysis. AccountingTools, referencia profesional consolidada para terminología contable y de control de gestión.

Corporate Finance Institute (2024). *EBITDA: Earnings Before Interest, Taxes, Depreciation and Amortization*. https://corporatefinanceinstitute.com/resources/accounting/what-is-ebitda/.

IFRS Foundation (2014). *IFRS 15: Revenue from Contracts with Customers*. Descuentos y bonificaciones se presentan como reducción de ingresos.

International Accounting Standards Committee (1983). *IAS 21: The Effects of Changes in Foreign Exchange Rates*. Revisión vigente: 2003. Tipo de cambio al cierre del período reportado para partidas monetarias históricas.

– (1992). *IAS 7: Statement of Cash Flows*. Revisión vigente: 2009. Recomienda el método directo para reporte gerencial de flujos.

Lumen Learning (2023). *Standard Cost Variance Analysis*. Open Educational Resource for Managerial Accounting. Distinción favorable/unfavorable según impacto en rentabilidad.