

EXECUTIVE INSIGHT · ENTERPRISE AI · SEPTEMBER 2026

From Data to Decision: the invisible layers of enterprise AI

Why AI stalls between pilot and production, and what leaders must build beneath the waterline: governed data, business meaning, prediction in the workflow and designed autonomy.

01 The iceberg: why AI stalls between pilot and production

02 Meaning: the layer nobody budgets for

03 The economics of meaning and the data-to-decision stack

04 The AI works. You decide: governed agentic applications

05 Vivix: industrial agentic AI in production, in Brazil

06 AI adoption is a leadership campaign

Insights from the S4 Digital × Siemens Digital Industries Software Executive Breakfast
São Paulo · 23 September 2026 · Gartner CIO & IT Executive Conference week
Published 28 September 2026

mendix Delivery Excellence Certification

Siemens · Mendix
Delivery Excellence Partner

SUMMARY

AI does not fail at the model. It fails at the foundation, in the meaning of the data and in governance.

OVERVIEW

On 23 September 2026, during the week of the Gartner CIO & IT Executive Conference, S4 Digital and Siemens Digital Industries Software brought a small group of decision-makers together in São Paulo. The morning had one question: how do you structure an AI journey that reaches results, from data to decision?

Six sessions followed: a Siemens view on trust and human capability, S4 Digital on the invisible layers of enterprise AI, a live Siemens build, Mendix in the AI era, the Vivix industrial journey and a keynote on AI strategy by author João Fornari Jr. The speakers came from different angles, yet they arrived at the same conclusion. This note consolidates what they said into a view for C-level leaders.

KEY FINDINGS

- **The bottleneck is the foundation, not the technology.** Speakers cited research showing 95% of organisations see no measurable P&L impact from generative AI.[1] Pilots run fast; production stalls.
- **Meaning is the missing layer.** Enterprises are not short of data. What they lack is the formalised meaning of that data: rules, exceptions and relationships. LLMs know the world. They do not know your company.
- **Ontologies and knowledge graphs have moved from theory to prerequisite.** They make answers traceable and cut cost, because agents stop rediscovering meaning at every query.
- **Value appears only when prediction reaches action,** inside the workflow, with a named person at each decision point.
- **Autonomy must be designed, not granted.** Governance is what allows AI to scale. It is not a brake.
- **It already works, in Brazil.** Vivix reported 85% faster resolution of production problems and 6,000 hours of manual work recovered in one year.[7]

RECOMMENDATIONS FOR C-LEVEL LEADERS

1. **Fund the foundation before the model.** Governed data, identity, lineage and access control come first. They are the part of the iceberg nobody budgets for.
2. **Treat business context as a proprietary asset.** Build an ontology and knowledge graph around one painful decision, then extend it. Models become commodities; your context does not.
3. **Put prediction inside the workflow.** Measure decision latency and cost per correct answer, not model accuracy.
4. **Write an explicit autonomy policy.** Decide what agents may execute, what needs approval and what stays with people, based on impact and reversibility.
5. **Judge AI on three tests, not one:** ROI, operational resilience and competitive advantage.[4]
6. **Run 90-day campaigns with mission teams.** Prepare the terrain, test in the real world, then decide: pivot, persevere or stop.[8]

STRATEGIC PLANNING ASSUMPTIONS CITED

> **By the end of 2026,** organisations will abandon 60% of AI projects that are not supported by AI-ready data.[3]

> **By the end of 2027,** over 40% of agentic AI projects will be cancelled due to rising costs, unclear value or weak risk controls.[3]

> **By 2028,** at least 15% of day-to-day work decisions will be made autonomously by agentic AI, up from 0% in 2024. 33% of enterprise applications will include agentic AI.[3]

> **By 2030,** 40% of application portfolios will be custom AI-native, up from 2% today.[3]

The morning at a glance. Six sessions, one thesis

SESSION	SPEAKER	CORE IDEA
AI and the return of human capability	Renata Sampaio, Siemens DISW	Turn intelligence into trust, and trust into action: data with context, AI with governance, technology tied to the business.
The invisible layers: from data to decision	Filipe Marques, S4 Digital	AI fails at the foundation, not the model. Meaning, captured in ontologies and knowledge graphs, is the missing layer.
Intelligence Center X in practice	Rafael Rocha, Siemens DISW	A governed agentic application built in days. The AI works; the expert decides at every critical step.
Mendix in the AI era	Nicolas Martinez, S4 Digital	Code generation is solved. Orchestration, integration and governance are the new bottleneck.
The Vivix AI journey	Aristóteles Terceiro Neto, Vivix	Industrial AI scales when the foundation is built layer by layer. More than 30 agents now run in production.
iA Arte da Guerra (keynote)	João C. Fornari Jr., Alper Seguros	AI adoption is a leadership campaign: risk of trust before ROI, the right battles, mission teams.

ANALYSIS

01 The iceberg: why AI stalls between pilot and production

95%
of organisations report no measurable P&L impact from GenAI
MIT[1]

63%
of breached organisations had no AI governance policy, or were still drafting one
IBM[2]

97%
of AI-related breaches involved inadequate access controls
IBM[2]

40%+
of agentic AI projects expected to be cancelled by end-2027
Gartner[3]

Almost everyone sees AI from the tip of the iceberg: the model, the assistant, the demo that impresses. What decides whether AI creates value sits below the waterline.

The numbers presented in São Paulo tell a consistent story. Investment is high. Returns are rare. The gap is not a shortage of algorithms. It is a shortage of foundations: data that is governed and carries context, controlled access, and a path from insight to action.

S4 Digital, drawing on McKinsey's State of AI survey,[5] framed the problem as three gaps. Most enterprises have at least two of them.

Gap 1: data without context

Data lakes store data, but they do not store meaning. "Customer" in the CRM, the ERP and billing are three different things, and none of the three systems knows it. AI cannot reason over systems that are not semantically connected. Recommendations stay generic, and nobody trusts them enough to act.

Gap 2: execution without intelligence

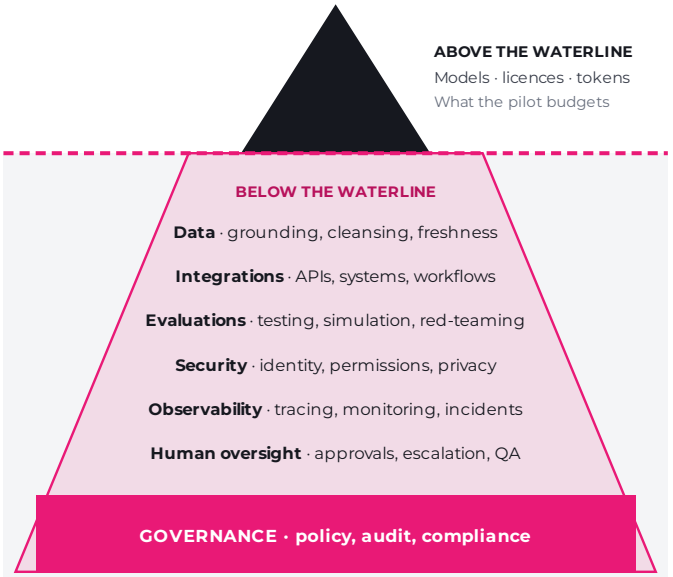
The model may be right, but the signal never reaches the system or the person where work happens. The gain shows up in the model metric and disappears in manual hand-offs. It is the classic dashboard that nobody uses to decide.

Gap 3: agents without governance

Shadow AI introduces unmanaged risk to security, data and software lifecycle. Black-box execution fails enterprise standards for traceability and audit. The pilot runs fast; production stalls.

> *AI does not fail at the model. It fails at the foundation and in governance.*
>
> FILIPE MARQUES, CEO LATAM, S4 DIGITAL

Figure 1. The AI cost iceberg: what gets budgeted, and what decides success



Ignored effort returns later as cost, delay or a pilot that never ships.

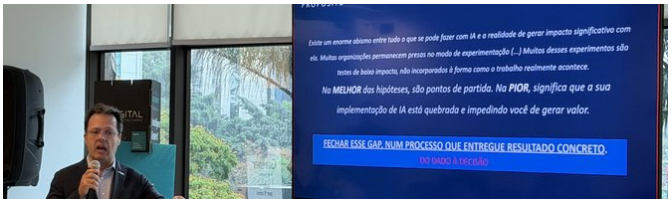
Source: S4 Digital, adapted from a widely shared practitioner analysis of AI costs, as presented on 23 September 2026.

THREE SIGNALS, ONE CONCLUSION

The builder. AI practitioners report costs rising as reasoning models turn one request into 20 to 50 intermediate steps. Most of this spend is still pilots. The real volume is ahead.

The CIO. Auana Mattar, CIO of TIM Brasil, reframed the old question of whether to build or buy. In the agent era, what matters is what you own: "proprietary data, experience and governance remain the hardest assets to copy."[6]

The ecosystem. At Dreamforce in September 2026, Siemens, Salesforce and NVIDIA announced joint work on agentic AI. Salesforce brings the customer and NVIDIA the compute. Siemens brings the operational reality that agents must reason about.



S4 Digital on closing the gap between what AI can do and the impact it delivers.



The audience: a restricted group of executives, board members, business owners and consultants.

ANALYSIS

02 Meaning: the layer nobody budgets for

A model can know what an invoice, a contract or a margin is. It cannot know what those things mean inside your company.

LLMs do not know your exceptions or your internal rules. They do not know why one customer is strategic, who may approve a given risk, or why two similar processes run differently in two companies of the same sector. That knowledge is highly idiosyncratic. It usually lives in people's heads, in spreadsheets and in code.

EXAMPLE · "FIFTY THOUSAND"

Renata Sampaio (Siemens) opened with a simple test: "50,000". Reais? People? Parts? Is it a lot or a little? For one customer a R$50,000 transaction is routine; for another it is far outside the pattern. The number is the same. Context changes its meaning. Source, period, unit, history and the decision it supports must reach the system. Otherwise the recommendation is wrong, however fluent it sounds.

EXAMPLE · THE SHIFT SUPERVISOR'S NOTEBOOK (1996)

In his first engineering project, Filipe Marques spent five months collecting production data at Autosil. At the time it was Europe's 4th-largest battery maker, running 24/7. The board already had almost all of that data, and it revealed nothing new. The breakthrough came from a shift supervisor's handwritten notebook. It recorded which mould sequences, machines and operators caused stoppages. Coded and analysed, it revealed the real process model. The fixes were simple: change mould sequences and standardise practices. **The data was always there. The meaning had never been formalised.** Today we would call it a taxonomy, an early ontology and a knowledge graph.

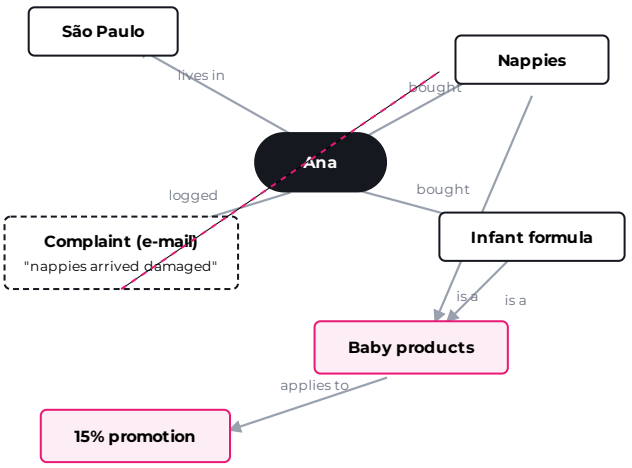
From taboo to trend

Until recently, "ontology" sounded academic, and proposing one could lose you the meeting. Two things changed that: the rise of Palantir Foundry, built on ontological modelling, and the arrival of ontology in Microsoft Fabric. In about two years it went from avoided topic to indispensable. The idea itself is old. Aristotle formalised categories of reality around 350 BC, and the semantic web and Schema.org followed. What is new is enterprise AI reasoning over internal data.



Renata Sampaio, Siemens DISW: "50,000. The value is the same. Context changes the meaning."

Figure 2. A database stores facts. A knowledge graph connects facts and gives meaning to the connections.



AGENTIC ACTION

Resolve the complaint first. Then offer the promotion with an apology, not a generic mailing.

Five structured systems and one unstructured e-mail resolve to the same entities, so the reasoning stays traceable rather than assumed. Source: S4 Digital, 23 September 2026.

Why this is a board issue

Without an ontology, business knowledge never exists explicitly anywhere. That prevents you from delegating important tasks and decisions to AI at scale. The model performs like a genius on one question and like a novice on the next, because it iterates over ambiguous data and tacit knowledge.

THE QUESTION FOR YOUR BOARD

If your favourite LLM could understand on its own everything that makes your business unique, what exactly would stop your competitors from reaching the same understanding?

The AI model can become a commodity. The context that makes your company unique never will. It should exist as its own layer: structured, governed, confidential and under your control.

> *Anyone can buy the model. Only you can build the graph of your business, and nobody can copy it.*
>
> FILIPE MARQUES, CEO LATAM, S4 DIGITAL



Filipe Marques, S4 Digital, on his 1996 Autosil project: the data was there, the meaning was not.

ARCHITECTURE

03 The economics of meaning and the data-to-decision stack

Ontology is not only about correctness. It is also about cost, because every token spent rediscovering your business shows up on the invoice.

Take the most ordinary question in any company: which customers have contracts expiring next month? It is already ambiguous. Is the "customer" the legal entity, the billing account or the parent group? Is "expiring" the end of term or the notice deadline?

Without shared definitions, each agent resolves the ambiguity with whatever context it finds. Two agents reach different answers, then spend more tokens reconciling them. Without a path through the data, they scan every document that mentions the customer. **You pay the machine to rebuild your company before it answers the question.**

With an ontology and a knowledge graph, every agent uses the same definitions and identifiers. The graph links entities to the source record, and data stays where it is. That means fewer retrieval calls, less irrelevant context and fewer attempts to reach the right answer.

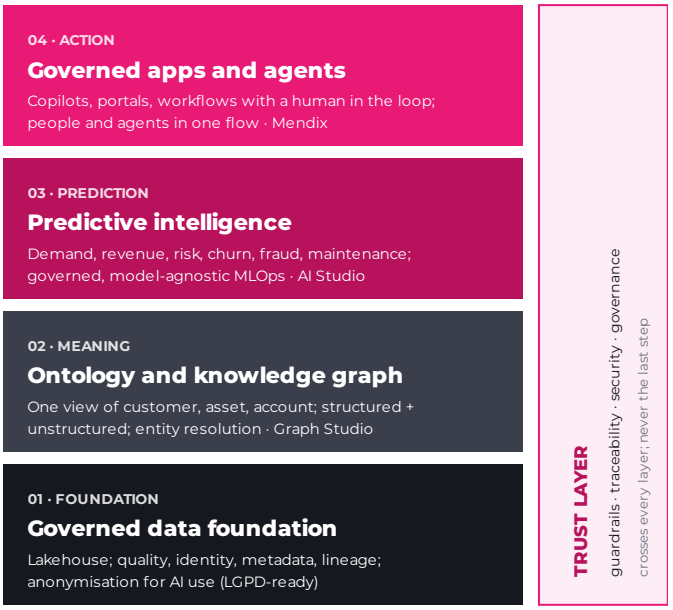
> **THE RIGHT METRIC**
>
> An ontology costs money to build and maintain, and it must pay for itself. The metric is not tokens per query. It is **cost per correct answer on repeated tasks, including maintenance.**

Knowledge graph versus context graph

A second distinction separates mature designs from expensive ones. The **knowledge graph** is what the company knows permanently: travellers, contracts, tariffs, policies, suppliers. It is stable, governed and has thousands of nodes. The **context graph** is what this decision needs, now. It is assembled at runtime and then discarded.

> **EXAMPLE · LUANDA, 03:12**
>
> A corporate traveller in Angola receives news of a family emergency. Six facts matter: the ticket is non-refundable; the next seat leaves at 06:40; the fastest route needs a transit visa he lacks; the client contract has a humanitarian exception; insurance covers repatriation; one named approver can authorise now. An agent given the whole warehouse would likely book the 06:40 flight and strand him at the connection. An agent given those twelve nodes solves the problem. **The difference is not the model. It is what the model was given.**

Figure 3. From data to decision: four layers and a trust layer that crosses all of them



Source: S4 Digital and Siemens DISW (Intelligence Center X), 23 September 2026.

The architecture presented by S4 Digital and Siemens follows the progression of this note. Data is made governed. Meaning is made explicit. Prediction is trained on data that carries meaning. Action happens where work is done. Siemens packages this as **Intelligence Center X**: knowledge graph, AI and ML modelling, agentic development and process orchestration, wrapped in an enterprise trust layer.

Two design choices stand out. First, **trust is not a module you add at the end.** Guardrails, traceability, security and governance cross all four layers. Second, **the stack is a ladder, not a package.** You can start with a knowledge graph that connects two systems, a model on data you already hold, or an application that acts on intelligence you already have.

Each layer is backed by a platform Gartner has recognised as a Leader. Mendix has been a Leader in enterprise low-code for nine consecutive years; AI Studio (listed as Altair) is a Leader in Data Science and ML platforms.[9]



Participants during the morning in São Paulo.

EXECUTION

04 The AI works. You decide: building governed agentic applications

Figure 4. SolutionMatch: nine steps from a request for proposal (RFP) to an explained bid. Every AI step sits between two human steps.

01 · YOU	02 · AI	03 · YOU	04 · YOU	05 · AI	06 · YOU	07 · AI	08 · AI	09 · AI
Ingest upload files as they arrived	**Extract** 4 agents build sourced reqs.	**Validate** accept, edit, reject, comment	**Context** choose knowledge graphs to use	**Match** norms, suppliers, past projects	**Decide** 3 ranked options with reasons	**Summarise** risks and gaps, export to PLM	**Predict** win probability, explained	**Learn** each outcome retrains the model

Human decision point: nothing moves on without it — AI works: multi-agent, graph-grounded, logged — Win probability: 51% → 84% → 87% as the expert adjusts

Source: Rafael Rocha, Siemens DISW, live demonstration on Intelligence Center X, 23 September 2026. Fictitious engineering-company scenario.

Siemens demonstrated a working application, SolutionMatch, rather than a concept. A solutions architect receives an RFP as PDFs, spreadsheets, e-mails and meeting transcripts. Four agents extract, review and reconcile the requirements, and each one stays traceable to its document, page and passage. The architect approves each one. **Human in the loop is not a principle on a slide. It is a button on the screen.**

Matching runs against knowledge graphs of ERP, manufacturing and supply-chain data. A vector search finds documents that mention "IP67". The graph finds a current catalogue product that meets IP67, with an active supplier, within cost and proven in a past project. That is multi-hop reasoning over relationships. The win-probability model, trained on projects the company has won and lost, shows which factors move the score. It lets the expert improve the bid before sending it, not after losing.

Built in days, not months

The application was built on the same platform it runs on. Maia turned one paragraph of briefing into epics, user stories and priorities, so the team's first meeting corrected a backlog instead of writing one. Concord ran Claude Code, Codex or Copilot agents inside Mendix Studio Pro to generate pages, workflows and microflows. The graph was exposed to agents over MCP, an open protocol. The prediction was published as an endpoint and called on demand.

> *The CIO's question is not "which AI should we use?" but "in which architecture will AI be embedded and governed?"*
>
> NICOLAS MARTINEZ, S4 DIGITAL

Figure 5. Prototype versus enterprise application

DIMENSION	AI PROTOTYPE	ENTERPRISE APPLICATION
Deliverable	Code, prompt, demo	Security, identity, trusted data, integration
Rules	Implicit in the prompt	Versioned, tested business rules
Operation	Outside production	Tests, observability, lifecycle
Scale	Isolated, hard to integrate	Common governed architecture

Source: Nicolas Martinez, S4 Digital, "Mendix in the AI era", 23 September 2026.

Figure 6. Not everything that can be delegated to AI should be

Autonomy	Approval	Human judgement
The agent may execute	A person authorises before the action	The decision stays with the person

LESS HUMAN CONTROL — MORE HUMAN CONTROL

Autonomy should follow impact, reversibility and policy. Source: Renata Sampaio, Siemens DISW, 23 September 2026.

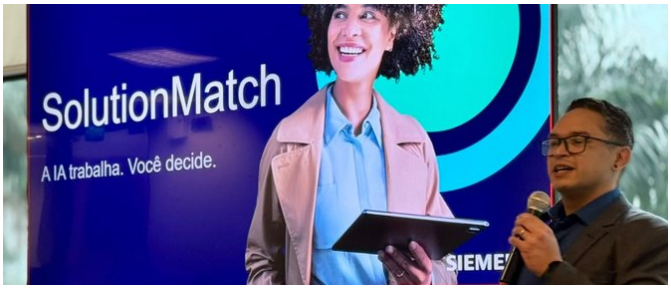
FOUR QUESTIONS BEFORE YOU TRUST AN ANSWER

Origin: where did the information come from? **Basis:** what supports the recommendation? **Limits:** what does this decision not cover? **Judgement:** where must a person step in? You must also be able to reconstruct what happened afterwards: what was recommended, what was done, who approved it and what the outcome was.

Governance is not a brake on AI. It is what lets it scale. Identity, human approval, guardrails, MCP security controls and lifecycle apply to agents exactly as they do to applications, and they are enforced centrally.



Rafael Rocha, Siemens DISW, explains the win-probability model: "not a magic number".



SolutionMatch, built on Intelligence Center X: "The AI works. You decide."

CASE STUDY

05 Vivix: industrial agentic AI in production, in Brazil

85%
reduction in time to resolve production problems

6,000 h
of manual work recovered in a single year

5d → <1
days to resolve a customer complaint

up to 4×
faster quality investigations

Vivix, Brazil's largest flat-glass maker, runs a R$1.5 billion plant producing 900 tonnes a day. It shows that the architecture in this note already runs in production, with sources any executive will recognise.

Aristóteles Terceiro Neto, who leads industrial transformation, was clear: **"Our transformation did not start with AI."** In 2021 Vivix had a strong operational-technology base in automation, robotics and computer vision, but it was disconnected from the business. In 2022–23 it connected shop floor, edge and cloud through a Unified Namespace and built its own industrial data fabric, with Mendix as its rapid development platform. Machine learning scaled in 2024, together with a structured start in generative AI through a **Virtual Engineer** for quality teams. In 2025 context reached the enterprise level. In 2026 agentic and physical AI are scaling across departments.

Today Vivix reports more than 32 industrial applications, 85% of the plant digitised and more than 30 agents and multi-agent systems in production. The Virtual Engineer combines Intelligence Center X with Claude on Amazon Bedrock and about 30 Mendix applications. It draws on SAP S/4HANA, Siemens Industrial Edge and Snowflake. Nothing was ripped out and replaced.[7]

ACCESSIBLE, CONTEXTUALISED, ACTIONABLE

Accessible: the right people and systems can actually use the data. **Contextualised:** the data is connected to the asset, the process and its business meaning. **Actionable:** it changes how the plant decides, operates and improves. "In our industry, data without context does not create intelligence. It creates noise."

LESSONS LEARNED

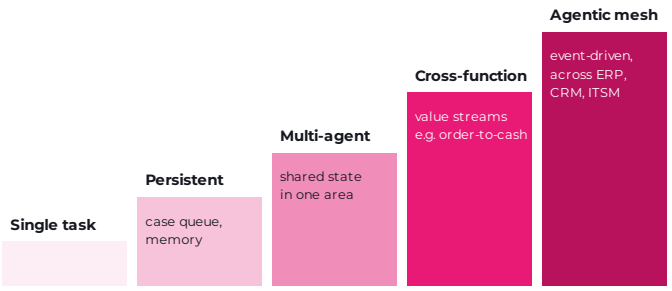
- Technology is not the hardest part. The main barriers are governance, business context and data semantics.
- Start with the business pain, not with AI tools.
- AI creates most value when it is embedded in daily industrial workflows.
- AI is not only incremental productivity. It redefines roles.

Figure 7. An AI maturity map, and the trap at each level

LEVEL	WHAT IT LOOKS LIKE	COMMON TRAP
N1 · Individual	People use AI to speed up their own tasks and decisions	"More tokens used = more AI-native"
N2 · Team	Copilots and agents enter team workflows	"More agents built = more AI-native"
N3 · Contextualised operation	Agents work on integrated context: data, systems, documents, assets	Connecting more sources ≠ more accurate results
N4 · Prescriptive	Agents recommend and act within defined limits and governance	Automating decisions without evidence or feedback quality
N5 · Adaptive	Systems learn from outcomes and adjust their own actions	Not yet observed at sufficient scale

Source: Aristóteles Terceiro Neto, Vivix Vidros Planos, 23 September 2026.

Figure 8. Five levels of agentic workflow: autonomy grows, and so does the need for governance



Source: Vivix, 23 September 2026.

The maturity map is one of the most useful tools from the morning. It warns against mistaking activity for progress. Token consumption and agent counts are signals of use, not value. The real step change is N3, when agents work on one shared industrial context. Even then, connecting more sources does not guarantee better answers unless meaning is governed.



Aristóteles Terceiro Neto presents the Vivix industrial AI journey.

LEADERSHIP

06 AI adoption is a leadership campaign, not a tool purchase

João Fornari Jr., author of *iA Arte da Guerra*, closed the morning with a reminder: technology gives access, strategy decides how it is used, and the organisation decides whether it scales.[8]

As APIs, models and cloud infrastructure become available to everyone, technology alone stops being an advantage. The advantage comes from the strategy for applying AI and from an organisation able to sustain it: data, culture, governance and vision.

Risk of trust before return on investment

Before scale, four pillars must stand: **fairness and bias**, continuously audited; **privacy and consent**, the social contract of data; **explainability**, since if nobody can explain a critical decision in business language the model is not ready for production; and **fiduciary responsibility**, which cannot be outsourced to the algorithm, the vendor or the black box.

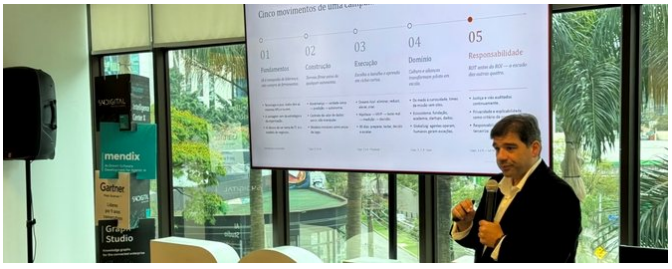
Choose the battle; prepare the terrain

Use AI not to do the same thing faster, but to make the competition irrelevant: eliminate long forms, cut resolution times from weeks to seconds, create predictive services. The right battle starts with real pain and ends with a value proposition that is hard to copy. And **do not start at the top**. Build in this order: governance, integration, intelligence, then autonomy within limits.

Move the troops

The biggest barrier is cultural. Leaders turn fear into curiosity through transparency and AI literacy. They break silos with OKRs that reward collaboration. They form **mission teams** that bring together domain, data, product, risk and operations around one real problem. **Governance is centralised; execution is decentralised.**

> *When AI gains the ability to act in the world, our responsibility for what we allow it to do grows too.*
>
> RENATA SAMPAIO, SIEMENS DIGITAL INDUSTRIES SOFTWARE



João C. Fornari Jr. on the five moves of an AI leadership campaign.

Figure 9. The 90-day plan: from intent to scale

DAYS 0–30	DAYS 31–60	DAYS 61–90
Prepare the terrain	**Test in the real world**	**Decide and scale**
Choose the battle	Launch the MVP	Pivot, persevere or stop
Audit risk of trust	Real data, real users	Prepare operations and training
Secure minimum data	Measure value and risk	Portability, monitoring
Form the mission team	Record the learning	

No owner, no plan. No metric, no experiment.
No trust limit, no responsible scale.

Source: João Fornari Jr., "iA Arte da Guerra", keynote, 23 September 2026.

MONDAY-MORNING CHECKLIST FOR THE C-SUITE

- Name one decision that takes days today and should take minutes.
- Map where its meaning lives: systems, documents and people's heads.
- Define "customer", "contract" or "asset" once, for everyone and every agent.
- Set the autonomy level for each action: execute, approve or judge.
- Measure cost per correct answer and time from data to decision.
- Report on three tests: ROI, resilience and competitive advantage.

A CLOSING QUESTION

If AI gave you back one hour of your day, what would you do with it?

Freed time does not turn into thought by itself. AI can give us back time. Giving that time a purpose is still our job.



Speakers at the close of the morning, with copies of *iA Arte da Guerra*.

SOURCES

Statistics reproduced as cited by the presenters on 23 September 2026. **1** MIT NANDA, The GenAI Divide: State of AI in Business 2025. **2** IBM, Cost of a Data Breach Report 2025. **3** Gartner forecasts and press releases, 2024–2026. **4** Gartner, "ROI Is Not Enough: The 3 Executive Tests of AI Value", 12 June 2026. **5** McKinsey, The State of AI, November 2025. **6** A. Mattar, MIT Sloan Management Review, 15 July 2026. **7** Siemens, Vivix customer case, June 2026. **8** J. C. Fornari Jr., iA Arte da Guerra. **9** Gartner Magic Quadrants for Enterprise Low-Code Application Platforms and for Data Science and ML Platforms, 2025.