

ORGANIZATIONAL DIGITAL TWINS: FROM SYSTEMS FRAGMENTATION TO COMPREHENSIVE GOVERNANCE

Alfonso Díez. Fundador and CEO

UGROUND Global SL · Spain

ABSTRACT

The ceramics industry operates with dozens or hundreds of applications, data repositories, and spreadsheets that coexist without an integrated view of the business. Digitalization has multiplied the amount of available data but has worsened fragmentation: each system speaks its own language, and the real knowledge of how the company operates resides in people, not in the technological architecture. This paper proposes an organizational digital twin architecture structured in three semantically distinct layers: the existing technology stack (L1), a single source of ontological truth (L2, SST), and a cross-cutting governance level with intelligent automation capability (L3). It argues that this three-layer separation is not a design choice but a structural necessity, describes ECMIA as a concrete implementation of this pattern, and illustrates it with the case of a distribution holding in the chemical industry that is applying this architecture to change its management model.

1. THE PROBLEM: THE CERAMICS COMPANY AS A FRAGMENTED SYSTEM OF SYSTEMS

A medium-sized ceramics company today operates as a system of systems: financial management ERPs, one or several MES systems, quality control platforms, maintenance tools, document management systems, production planning applications, internal communication platforms, and, as an invisible yet omnipresent substrate, dozens or hundreds of spreadsheets that cover the gaps no application covers well. Added to this are the systems of suppliers, customers, and logistics operators with which the company exchanges information in a more or less structured way. The literature on IIoT, digital transformation in production, and smart manufacturing shows that this proliferation of devices, applications, and digital services increases data capture capacity, but also raises the complexity of integrating, coordinating, and governing information [12][13][14].

The result is not an integrated system. It is a set of coexisting systems, each with its own data model, its own semantics, and its own update cycle. A customer order can be one entity in the ERP, another in the MES, another in the quality system, and another in the Excel spreadsheet the production manager uses to plan the week. These are not the same entity represented in four different ways: they are partial, misaligned, and often inconsistent representations of an operational reality that no system captures completely. This type of heterogeneity is precisely one of the problems that the literature on knowledge graphs and semantic technologies identifies as a structural condition of digitalized industrial environments [15][16].

This situation has direct consequences on decision-making capacity. When a problem arises in the supply chain, when a regulatory change must be addressed, or when management needs a consolidated view of operational performance, the usual response is to convene a meeting, gather data from multiple sources, and manually build a view that does not exist automatically. The knowledge of how the company actually operates resides in people, not in the technological architecture. If those people leave, the knowledge is lost, degraded, or limited to informal procedures. Studies on knowledge graphs in manufacturing point out precisely that knowledge fusion, semantic contextualization, and traceability of entities and relationships are recurring needs for progress toward a more flexible and governable Industry 4.0 [16].

This is not a new problem, nor one exclusive to the ceramics sector. What is specific to the current moment is the scale of the problem and the urgency of solving it. Global competitive pressure, supply chain volatility, traceability and sustainability requirements, and the rise of artificial intelligence as an operational tool have turned systems fragmentation into a strategic constraint. Recent data governance frameworks and reports on enterprise AI adoption insist that the ability to generate value with analytics and artificial intelligence depends on having governed, consistent, and usable data at scale [1][2]. In the field of logistics and supply chain, recent literature on digital twins likewise underscores the need for integrated, transparent, and resilient models to support decisions in dynamic, multi-actor environments [17][18]. Companies unable to build an integrated view of their operations will face growing difficulties competing in environments that demand fast, traceable, and well-founded responses.

2. WHY CURRENT APPROACHES FAIL

Faced with this diagnosis, organizations have tried different responses for years. The most common has been point-to-point integration: connecting systems through APIs, middleware, or custom development so they can share data. This approach improves operational synchronization but does not solve the underlying problem: each integration requires mapping two different data models, with specific transformation rules and technical dependencies that grow cumulatively. This integration debt appears recurrently in industry diagnostics on enterprise connectivity [3]. Furthermore, the literature on digital transformation in production and on IIoT shows that connectivity is a necessary but not sufficient condition when the connected systems still retain local semantic models and there is no shared representation of the business [13][14].

A second response has been the implementation of Business Intelligence platforms, corporate data warehouses, or data lakes. These solutions are useful at the analytical level, but they generally operate on historical or consolidated data for observation, not on ongoing operational reality. In other words, they help see what has happened but are not enough to govern cross-cutting processes, execute corporate policies, or act in a coordinated way across multiple source systems. The literature on knowledge graphs draws a precise distinction between analytical data repositories and semantic models capable of representing entities, relationships, context, and interpretation rules, which enables reasoning over heterogeneous information rather than merely storing it [15][16].

The third response, increasingly common, is the adoption of artificial intelligence solutions: copilots, conversational assistants, predictive models applied to specific processes, or recommendation mechanisms. These tools can generate real value in bounded use cases, but they amplify the structural problem when deployed on a fragmented architecture. The quality of AI depends directly on the quality, consistency, and traceability of the data that feeds it; therefore, without data governance there is no trustworthy AI [1][4]. In digital twin approaches, AI is not conceived as a substitute for the knowledge model, but as a capability that operates on an explicit semantic representation of the domain, able to contextualize entities, processes, relationships, and decision alternatives [19].

A metaphor can help us better understand this situation. We can picture the company shaped like the letter π . The vertical lines are the local silo applications (ERP, CRM, etc.) that resolve atomic functions (issuing an invoice, paying payroll, registering an order, etc.). The horizontal line represents the integrated or corporate functions, which connect the vertical functions and work on second-order processes (optimizing cash flow, increasing market penetration, redesigning logistics flows, etc.). Second-order processes have their own internal logic and semantics; they are not merely an interconnection of local functions, nor can they be resolved with AI agents. Not only that, but they are key to the company's success in resilience and competitive capacity.

3. THE THREE-LAYER SEPARATION IS SEMANTICALLY NECESSARY

The proposal to structure enterprise knowledge architecture in three layers is a consequence of the semantic nature of the problem: each layer represents a distinct type of knowledge, with radically different properties, scope, and technical requirements. This is a way of digitally transforming an organization through a Model-Based Systems Engineering (MBSE) approach, establishing an explicit separation between the operational model (what the system does), the

semantic model (what the system knows about itself), and the governance model (the rules that determine how the system should behave). This distinction, developed in the context of complex systems engineering, is directly applicable to the problem of enterprise architecture [5].

Another source of inspiration for this modelling approach is SoSE (Systems of Systems Engineering). Within this framework, we understand the organization not as a deterministic machine but as a complex adaptive system, whose different subsystems evolve in search of various equilibrium patterns. The three-layer organization makes it possible to preserve the operational independence of the subsystems while maintaining coordination and balance at the level of the overall system [10].

3.1. Layer L1: Local Domain Models

The existing technology stack (ERPs, MES, quality systems, CRMs, spreadsheets, documents, etc.) defines a semantics of specific scope. Each system correctly represents the reality of its functional domain but does not attempt, nor can it, represent the company's integrated view. A warehouse management system correctly models locations, references, and movements, but knows nothing about how that stock affects commercial commitments, production planning, or financial indicators. This limitation is not a design flaw but a natural consequence of functional specialization.

The frequent strategic mistake consists of trying to extend these local models to cover cross-cutting needs, whether through ERP customizations, custom development, or the accumulation of point-to-point integrations. The result is systematically the same: overloaded, fragile, and costly-to-maintain systems that still fail to solve the underlying problem.

3.2. Layer L2: The Single Source of Ontological Truth (SST)

No L1 model attempts to represent the ontologically integrated view of the company, because that is not its function. Building that view requires an intermediate layer specifically designed for it: the Single Source of Truth (SST). This layer does not replicate the data from the source systems; it acts as a semantic layer over them, integrating partial representations into a coherent model where each business entity, customer, order, supplier, asset, process, exists only once with its complete meaning and traced relationships. The design model (the ontology) that implements the SST is corporate, not departmental, and specifically incorporates data governance features such as origin, explainability, status, version, and change control.

The SST requires a specific technical infrastructure that is not equivalent to a conventional relational database or a data warehouse. It needs to be agnostic with respect to the source data models (SQL, NoSQL, files, APIs), capable of representing complex semantic relationships, and, in the current context, vectorized to enable semantic search and reasoning over content through generative AI techniques. This type of repository is referred to in the knowledge management literature as a knowledge graph enriched with semantic capabilities [15][16].


3.3. Layer L3: Cross-Cutting Governance and Intelligent Automation

L3 represents the category of enterprise knowledge that does not belong to any L1 system and that also cannot be resolved solely through a better representation of data. It is the cross-cutting layer of π . For example: (1) a credit policy that simultaneously affects sales, finance, and logistics; (2) a quality criterion that determines both the acceptance of raw materials and the classification of the finished product; (3) a response procedure for a stock-out that involves purchasing, production, and customer service.

These rules exist in no system because a) they fall outside the application domain of any individual system, or b) they would require modifying multiple systems at a cost, risk, and maintenance complexity that no organization is willing to accept. The result is that they reside in documented procedures, in people's tacit knowledge, or simply are not applied consistently.

Layer L3 is the space where these cross-cutting rules can be modelled, applied, and evolved without touching the source systems. Its conception is comprehensive, not vertical: it can simultaneously manage information and business rules covering any aspect of the organization. This is where the intelligent automation, decision support, and cross-cutting process management capability that characterizes a mature organizational digital twin resides.

Layer L3 supports stakeholders: employees, executives, customers, suppliers, partners, public administrations, regulators, etc. This support is structured around three main services: **decision support** (queries, dashboards, data analytics, and RAG-based natural language search); **autonomous agents** capable of making intelligent decisions about the business and events and of generating data flows toward L2 and L1; and management of **business processes** that are cross-cutting, based on cases identified in L2.



4. A THREE-LAYER ARCHITECTURE AS A GOVERNANCE PATTERN: ECMIA AS A REFERENCE IMPLEMENTATION

The three-layer separation described in the previous section is not a theoretical proposal without practical grounding but an architectural pattern that can be materialized in different implementations. What determines the viability of any specific implementation is the ability to maintain semantic coherence across the three layers: that L1 remains the operational system of record; that L2 is genuinely a corporate ontological representation, not a renamed data warehouse; and that L3 operates on externalized business rules, not on logic written into applications.

ECMIA (Enterprise Control & Management Intelligent Architecture) is an implementation of this pattern developed by UGROUND. Its three-layer structure corresponds exactly to the semantic separation described. The component that articulates the three layers is ROSE, an ontological engine based on recursive semantic engineering [7] that acts as a digital twin. ROSE is not one of the layers: it manages the acquisition of information from L1 into the SST, the construction and maintenance of the ontological model in L2, and the exploitation of that model in L3 through three main applications: Decision Support, Process Management, and Agentic Automation.

A relevant property of this architecture is its non-invasive nature: by deploying as a semantic layer over existing systems without replacing or modifying them, adoption cost is significantly reduced compared to platform replacement initiatives, and the first use cases can be operational within a relatively short timeframe. This property is a direct consequence of the separation between layers: L1 remains unaltered because corporate semantics resides in L2, not in the source systems.

4.1. The Role of AI in the Architecture

A relevant aspect of this architecture is the conception of the role of artificial intelligence. In many digital transformation projects, AI is presented as the product, the copilot, the assistant, the agent, the solution. In a three-layer architecture like the one described, AI is a tool that operates at multiple points; it is not the central axis.

In the acquisition layer (L1→L2), AI is applied to enrich and classify information: entity extraction, resolution of semantic inconsistencies, linking of objects across source systems. In the SST (L2), semantic vectorization enables search and reasoning over the repository's content. In the exploitation layer (L3), AI intervenes in all three applications: generating analysis and recommendations in Decision Support and allowing Ontology Exploitation using natural language in user interfaces, orchestrating flows in Process Management, and executing autonomous tasks within the limits defined by business rules in Agentic Automation.

This distribution has an important practical consequence: the quality of AI depends directly on the quality of the ontological model. A recommendation system that operates on a well-built SST with an explicit knowledge graph produces qualitatively superior results compared to one that operates on fragmented and semantically inconsistent data.

5. AN EXAMPLE OF APPLICATION IN A CORPORATE HOLDING

To illustrate the practical value of the three-layer architecture, we present a real-world case of application to a corporate holding. This holding (chemical distribution sector) is fragmented into more than 20 companies operating across four connected business lines. Although they have functional independence, they share many synergies in suppliers, products, warehouses, customers, finance, and procedures.

5.1. The Problem in Terms of Fragmentation

Its L1 architecture has grown from an original ERP that has been cloned over time. This technology has enabled the group's growth but has created a growing burden on the ability to manage the group as an integral system. The result is very common: spreadsheets have multiplied, managerial decision-making capacity is limited by data access and quality, certain key processes lack adequate infrastructure, and synergies between companies cannot be leveraged. Faced with this reality, the company has two possible strategies: consider replacing the ERP and adding several additional vertical systems or implement a cross-cutting corporate digital twin that builds on what already exists and generates all the additional value the company needs.

5.2. The Solution with the Three-Layer Pattern

Deployed onto this reality, the three-layer architecture builds a unified representation in the SST (L2) from multiple sources of information: on one hand all the ERPs, and additionally Excel

spreadsheets, PDF documents, and a new interconnection with banks. All of these sources provide information on suppliers, purchases, customers, sales, banks, products, stock, prices, etc. All of them are entities related within a single ontological model, updated in near-real time from the source systems without modifying them.

Building on this SST, layer L3 takes charge of the key cross-cutting processes that fall outside the scope of L1. The processes considered most important in the first development cycle are:

- a) the managerial decision-making model, for the group and for each of the companies;
- b) warehouse optimization, integrating all information to obtain group synergies;
- c) the treasury model integrating banks, cash flow forecasts, and simulators;
- d) budget-based management; and
- e) a new purchasing model that integrates the group's procurement processes along with its warehouses and stock.

In a process of this type, timing and development dynamics are key. These need to be very fast processes with an iterative design cycle: moving from the design board to MVP-mode testing, followed by ongoing refinement in continuous cycles. It is not traditional waterfall development, but neither is it an agile methodology; rather, it is an iterative refinement model in which the final design gains richness and nuance with each cycle.

The timing applied to this project is six months end-to-end, with a cycle of four iterative deployments to generate business value while the investments are being made.

6. VALIDATION: FROM THEORY TO INDUSTRIAL PRACTICE

The viability of ontology-based organizational digital twins is not a hypothesis. It is validated both by large-scale international experiences and by applications in the European market. Internationally, companies such as Microsoft, Google, and SAP have announced new products based on ontological developments. Another company, Palantir Technologies, constitutes a relevant reference for validating the scalability of the ontological pattern in highly complex environments. Its platform has addressed radically different problems without changing the underlying architecture: from optimizing aircraft fleet maintenance to the operational management of hospitals [6].

In the European market, UGROUND provides evidence of the pattern's applicability across organizational environments of diverse nature. With more than fifteen years of deployments in demanding environments, its track record spans sectors as diverse as banking, insurance, industry, distribution, food, healthcare, staffing, telecommunications, and parcel delivery, in both the private and public sectors. This sectoral diversity is in itself evidence of the pattern's cross-cutting nature: the same digital twin architecture solves knowledge governance problems in organizations with completely different structures, processes, and source systems.

7. STRUCTURAL PROPERTIES OF THE PATTERN AND IMPLEMENTATION CONDITIONS

The three-layer architecture presents a set of structural properties that distinguish it from conventional integration approaches, precisely because each layer does what it is meant to do, and only that. The most relevant properties are:

Preservation of the existing technology stack. The accumulated technology investment is preserved and enhanced. There is no system replacement, data migration, or operational implementation risk arising from a platform change.

Non-invasive nature. By operating as a semantic layer over existing systems, the first use cases can be operational within a relatively short timeframe. The model is built incrementally.

Technological agnosticism. The architecture integrates with any source system, regardless of its technology, age, or vendor. It creates no dependency on a specific technology ecosystem.

Adaptability in high-uncertainty environments. The ontological model can evolve without reprogramming systems. When business rules change, the model in L3 or the ontology in L2 is updated, not the application code.

Experience suggests that the critical success factor is not technological but organizational: the company's willingness to undertake a process of integrating its management model. This process of capturing tacit organizational knowledge is, in itself, one of the most valuable assets generated by the implementation, regardless of the direct operational benefits.

8. CONCLUSION: FROM DATA TO BUSINESS GOVERNANCE

The ceramics industry has spent years accumulating digital capabilities: sensors, MES, analytics, artificial intelligence applied to specific processes. The question that still lacks a satisfactory answer in most organizations is not how to digitalize more, but how to make everything that has been digitalized work as a coherent system.

The proposal of this paper starts from a structural diagnosis: systems fragmentation is not solved with more integrations or better BI tools. It is solved by recognizing that the three types of enterprise knowledge – operational domain knowledge (L1), integrated ontological knowledge (L2), and corporate governance knowledge (L3) – have different natures and require different architectural layers. ECMIA is a concrete implementation of this pattern; the pattern itself has validity independent of any specific implementation.

This separation is not a proposal for added complexity. If each layer does what it is meant to do, and only that, then the whole becomes simpler, more maintainable, and more evolvable. The model can be updated as the business changes, artificial intelligence can be applied where it generates real value, on data with meaning and traced relationships, and decisions can be made with a complete view, not with fragments.

This approach has a critical second-order effect for companies: maturity in their management model. This model moves from being implicit (spreadsheets, meetings) to explicit (management layers, analytical systems, integrated information, corporate dashboards, etc.), becoming a key factor in the evolution of companies' strategy.

9. REFERENCES

- [1] Gartner. Data and Analytics Governance Framework. 2025.
<https://www.gartner.com/en/data-analytics/topics/data-and-analytics-governance>
- [2] McKinsey & Company. The State of AI in 2025. McKinsey Global Institute, 2025.
- [3] MuleSoft / Salesforce. Connectivity Benchmark Report 2024.
<https://www.mulesoft.com/lp/reports/connectivity-benchmark>
- [4] [Huyen, C. AI Engineering. O'Reilly Media, 2025.
<https://www.oreilly.com/library/view/ai-engineering/9781098166298/>
- [5] Holt, J.; Perry, S. Implementing MBSE into Your Business – The Trinity Approach. INCOSE UK, 2020.
- [6] Palantir Technologies. Ontology-Driven Operations: Case Studies in Aerospace, Healthcare and Defense. Palantir White Paper, 2024.
<https://www.palantir.com/platforms/foundry/>
- [7] [UGROUND Global SL. Recursive Ontology-Based Systems Engineering. US Patent US9760345B2. USPTO, 2017. <https://patents.google.com/patent/US9760345B2>
- [8] DAMA International. The DAMA Guide to the Data Management Body of Knowledge (DAMA-DMBOK2). Technics Publications, 2017.
- [9] INCOSE. Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities, 4th ed. Wiley, 2015.
- [10] Boyes, H.; Hallaq, B.; Cunningham, J.; Watson, T. The industrial internet of things (IIoT): An analysis framework. Computers in Industry, 101, 1–12, 2018.
- [11] Brauner, P.; Dalibor, M.; Jarke, M.; Kunze, I.; Koren, I. A Computer Science Perspective on Digital Transformation in Production. ACM Transactions on Internet of Things, 3(2), Article 15, 2022.
- [12] Hogan, A. et al. Knowledge Graphs. ACM Computing Surveys, 54(4), Article 71, 2021.
- [13] Buchgeher, G.; Gabauer, D.; Martinez-Gil, J.; Ehrlinger, L. Knowledge Graphs in Manufacturing and Production: A Systematic Literature Review. IEEE Access, 9, 55537–55554, 2021.
- [14] Le, T. V.; Fan, R. Digital Twins for Logistics and Supply Chain Systems. Transportation Research Part E, 179, 103315, 2023.
- [15] Zhang, J. et al. Supply Chain Digital Twin Framework Design. Proceedings IEEE IEEM, 2021.
- [16] Rožanec, J. M. et al. Actionable Cognitive Twins for Decision Making in Manufacturing. International Journal of Production Research, 60(2), 452–478, 2022.
- [17] Karabulut, E.; Pileggi, S. F.; Groth, P.; Degeler, V. Ontologies in Digital Twins: A Systematic Literature Review. Future Internet, 15(10), 335, 2023.