Research Article

A Governable Digital-Twin Framework for Transnational Manufacturing Networks

Lucia Rodríguez García [1,*]
[1] Department of Computer Science, Autonomous University of Madrid , Spain
* Corresponding author: Lucia Rodríguez García
Email: LuciaRodríguez@outlook.com

***Abstract*:** Current research on digital twins provides a solid foundation for asset synchronization and cross-border supply chains, but it fails to address issues such as data quality management across organizational boundaries, model uncertainty, network risks, and decision-making authority. This study constructs a trustworthy network digital twin architecture, differentiating perception and event capture, edge verification, semantic interoperability, federated twin services, and human-machine decision coordination through key findings from integrated standards and peer-reviewed research. Traceability, uncertainty management, and policy enforcement are integrated across these layers. This framework aims to promote autonomy and network coordination at the local plant level, but sensor drift, semantic heterogeneity, logistical delays, and jurisdictional limitations may restrict its portability. Further research is needed through simulation, hardware-in-the-loop experiments, and multi-site comparative field studies.

Keywords: *Digital twin, Cyber-physical manufacturing, Supply-chain resilience, Semantic interoperability, Edge computing, Trustworthy AI.*

1. Introduction

Manufacturing networks are no longer confined to a single plant. Production schedules, quality decisions, inventory positions, transport capacity, energy constraints, and compliance obligations interact across firms and jurisdictions. This interdependence creates a coordination problem: operational decisions must be made quickly at the shop floor while remaining consistent with slower network-level planning and governance processes. Research on enterprise-wide informatization architectures and standards-driven management systems suggests that technical integration may require an institutional architecture extending beyond isolated software projects [1][2]. These studies are relevant to governance design, although their practice-oriented methods do not, by themselves, establish the real-time fidelity expected of a manufacturing digital twin.

The literature nevertheless leaves an important gap between asset-level twins and network-level decision systems. Kritzinger et al. [3] used a categorical review to distinguish digital models, digital shadows, and digital twins, a method that remains valuable for preventing conceptual inflation, although its manufacturing-centered classification offers limited guidance for decisions spanning independent firms. Kamble et al. [4] synthesized a broad supply-chain literature and proposed an implementation framework, yet the heterogeneity of the reviewed studies makes it difficult to infer which architectural components produce particular sustainability outcomes. Ivanov and Dolgui [5], working through disruption-oriented modeling, showed how a supply-chain twin may support resilience analysis; their abstraction is highly relevant to network coordination, while the organizational and semantic costs of maintaining real-time cross-company data are less directly resolved. These methodological differences suggest possible complementarity rather than a single settled design.

This article does not claim that a unified architecture has already produced quantified industrial benefits. Instead, it asks: How should a transnational manufacturing-network digital twin be structured so that its recommendations remain traceable, uncertainty-aware, secure, and governable? The article makes three contributions. First, it defines a five-layer architecture that connects plant-floor cyber-physical systems with network planning. Second, it introduces cross-cutting assurance controls for provenance, uncertainty, and policy. Third, it specifies testable propositions and a staged validation agenda that separates technical performance from operational and organizational outcomes.

2. Scope and Review Method

The paper uses a focused conceptual synthesis rather than a systematic review or an empirical case study. The source base was selected to cover four constructs required by the proposed framework: manufacturing digital-twin reference architectures, cyber-physical production systems, supply-chain twins and resilience, and edge-enabled analytics. Standards were used to define architecture and interoperability requirements, whereas peer-reviewed reviews and foundational studies were examined for their modeling assumptions, units of analysis, evidence types, and reported limitations. Because the search was purposive rather than exhaustive, relevant sector-specific evidence may have been omitted; the synthesis should consequently be read as theory-building work whose adequacy depends on later comparison with broader systematic reviews.

The unit of analysis is a manufacturing network comprising multiple plants and external logistics nodes. The framework includes production and material-flow decisions, but it does not attempt to model every enterprise function. Occupational or biometric monitoring is excluded because it requires distinct ethical, medical, and privacy protocols and was not supported by the evidence base of the original manuscript. Environmental performance can be represented when auditable measurements and system boundaries are available, but the framework does not infer carbon or water savings from operational proxies alone.

2.1 Research exploration and scope adjustment

The conceptual development was not entirely linear. An early attempt treated compliance documents, equipment signals, logistics events, environmental indicators, and worker-related information as comparable streams within one architecture. During source comparison, this aggregation proved difficult to defend: standards described manufacturing elements at different levels of granularity, review articles used partly overlapping definitions of a twin, and occupational-health information introduced ethical and inferential questions that could not be addressed merely through an additional data layer. The scope was

consequently narrowed to production and material-flow coordination, while environmental variables were retained only where measurement boundaries could be made explicit. A second adjustment followed from the tension between flexible schema-on-read ingestion and the stable semantics required for operational decisions. Exploratory ingestion was retained as a possible discovery mechanism, but validated, version-controlled mappings were made mandatory before decision use. These changes reduced the apparent breadth of the framework, although they may improve its conceptual defensibility.

3. Design Requirements

3.1 Purpose-bounded fidelity

A twin should be as detailed as its decision purpose requires, not a maximal replica of the entire network. Asset-health decisions may require high-frequency vibration or temperature data, whereas allocation decisions may rely on hourly production capacity and transport events. Each model therefore needs an explicit purpose, owner, update rate, validity range, and retirement criterion. This requirement limits unnecessary data collection and reduces the risk that a model is reused outside the conditions for which it was validated.

3.2 Semantic and temporal interoperability

Cross-site coordination fails when identical labels describe different quantities or when timestamps, units, and aggregation intervals are inconsistent. Enterprise "Standard+ System" approaches may offer an organizational analogy for combining stable rules with adaptable applications [2], while cyber-physical manufacturing research more directly emphasizes the connection between physical processes and computational decision mechanisms [6][7]. The proposed architecture consequently requires canonical identifiers, unit metadata, event-time semantics, versioned mappings, and explicit treatment of late or missing events, although the cost of maintaining these mappings across firms remains difficult to estimate.

3.3 Local autonomy with network coordination

Edge computing can reduce latency and preserve operation during intermittent connectivity [8]. More specialized scheduling studies have considered fault tolerance and latency-aware allocation for edge-AI workloads [9][10], yet their reported architectures may not transfer directly to heterogeneous factory equipment or cross-border communications. Low latency alone does not justify autonomous control. Local controllers should retain authority over safety-critical and time-critical actions, while the network twin coordinates decisions such as order allocation, inventory positioning, and disruption response. Conflicts between local and network objectives must be resolved through predefined escalation rules rather than hidden optimization weights.

3.4 Trustworthiness by construction

Manufacturing twins expand the attack surface because they connect models, operational technology, and enterprise systems. Digital-twin-based anomaly detection illustrates how physical and model-based evidence can support cyber-attack detection, while also showing the need to distinguish malicious behavior from process anomalies [11]. Every recommendation should therefore be linked to its input data, model version, uncertainty estimate, policy checks, and approving actor.

3.5 Organizational adoption and industrial ecosystem conditions

Technical architecture alone may not explain adoption. Research on the commercialization of AI-security technologies and the transformation of enterprise informatization achievements emphasizes investment pathways, industry–university–research–user coordination, and the movement from technical prototypes toward market use [3][4]. Work on strategic materials parks and large-scale innovation funds further suggests that shared infrastructure and patient financing could support ecosystem formation [12][13], although evidence from a particular park or funding program cannot be assumed to generalize to transnational manufacturing networks. These contributions are used here to identify institutional variables rather than to validate the proposed architecture.

Capability development is another possible constraint. A teacher-training program for AI technology [14] concerns an educational setting rather than industrial operations, but its evaluation-oriented approach draws attention to the fact that user competence must be assessed rather than presumed. In a manufacturing context, training should cover uncertainty interpretation, escalation procedures, cybersecurity hygiene, and the reasons an operator may override a recommendation. The relevance is indirect, and further research is needed to determine whether educational-program evaluation methods can be adapted to industrial teams.

The implementation problem is likely to be especially uneven among small and medium-sized enterprises. Studies of multi-format management systems and long-term digital-transformation mechanisms for smaller firms[15][16] describe organizational continuity and scalable management as central concerns. Their broad managerial orientation does not resolve model fidelity or cyber-physical timing, yet it raises a significant boundary condition: a network twin designed around the resources of large manufacturers may shift data-maintenance burdens onto smaller suppliers, thereby weakening participation or creating systematic blind spots.

4. Proposed Architecture

The architecture contains five functional layers and three assurance planes. The layers define where transformations occur; the assurance planes define constraints that apply throughout the system.

Layer	Function	Representative components	Required output
L1	Sensing and event capture	Machine signals, quality events, inventory movements, transport milestones	Timestamped observations with source identity
L2	Edge validation and feature extraction	Unit checks, plausibility tests, drift indicators, local feature computation	Quality-scored data and local state estimates
L3	Semantic interoperability	Canonical identifiers, unit conversion, ontology mappings, event-time alignment	Versioned, context-rich network events
L4	Federated twin and analytics	Plant twins, logistics-node twins, scenario models, uncertainty-aware	Predictions, scenarios, and confidence bounds

optimization

L5	Decision orchestration	Policy engine, human approval, local execution interfaces, audit log	Authorized action or documented rejection

4.1 Layer 1: Sensing and event capture

Layer 1 records observations without assuming that every observation is correct. Source identity, calibration status, sampling rate, unit, and event time travel with the measurement. Logistics events must similarly identify the reporting party and distinguish planned, estimated, and confirmed milestones.

4.2 Layer 2: Edge validation and feature extraction

Layer 2 performs bounded local processing: range and rate-of-change checks, unit validation, duplicate detection, signal-quality assessment, and features required for time-critical control. Suspect data are flagged rather than silently imputed. When connectivity fails, local systems continue within an approved operating envelope and later reconcile events with the network twin.

4.3 Layer 3: Semantic interoperability

Layer 3 maps local tags and business events to versioned network concepts. Mappings are treated as governed software artifacts with owners, tests, and change histories. Schema-on-read may be useful for exploration, but production decisions require validated mappings and explicit handling of unknown fields. This distinction prevents flexible ingestion from becoming uncontrolled semantic drift.

4.4 Layer 4: Federated twin and analytics

Layer 4 maintains separate but connected twins for plants, assets, products, and logistics nodes. Federation avoids a single monolithic model and permits jurisdiction-specific data retention. Scenario services compare feasible actions under disruption, while prediction services report uncertainty and out-of-distribution warnings. Optimization outputs are recommendations until they pass policy and authority checks.

4.5 Layer 5: Decision orchestration

Layer 5 converts analytical output into accountable decisions. A policy engine checks safety constraints, contractual commitments, access rights, and jurisdictional restrictions. Human authorization is mandatory for decisions with material safety, workforce, regulatory, or cross-company consequences. Automated execution is restricted to pre-approved, reversible actions inside a validated operating envelope.

5. Cross-Cutting Assurance Planes

5.1 Provenance and auditability

Every state estimate and recommendation should be reproducible from retained inputs, transformation rules, model versions, and policy decisions. Provenance must cover both data and semantics: a value is not auditable if its numerical origin is known but the meaning of its field changed between sites or software versions.

5.2 Uncertainty and model risk

The system should represent measurement uncertainty, model uncertainty, and scenario uncertainty separately. Confidence bounds must not be interpreted as guarantees. Drift detection, out-of-distribution testing, and site-specific recalibration are required before models are transferred between plants. Missing data should trigger a declared fallback mode rather than invisible replacement.

5.3 Security, privacy, and policy enforcement

Identity, least-privilege access, network segmentation, signed model artifacts, and incident logging are baseline requirements. Data-sharing rules should be enforced close to the data source, with only the minimum features required for a network decision exposed. Policy changes must be versioned and tested because a correct model can still produce an impermissible action.

6. Testable Propositions and Validation Agenda

P1 — Semantic quality: Versioned mappings and explicit event-time semantics may reduce cross-site reconciliation errors relative to ad hoc tag-level integration, although additional governance work could offset part of the operational benefit.

P2 — Operational continuity: Edge validation with approved fallback envelopes is expected to reduce decision latency and service interruption during network degradation, but the effect may depend on whether local models remain valid during the disruption.

P3 — Decision quality: Uncertainty-aware scenario ranking may produce fewer infeasible or unstable recommendations than point-estimate optimization under sensor drift and delayed logistics events; overly conservative uncertainty estimates could, however, suppress useful actions.

P4 — Governance: Provenance-linked human authorization may improve post-event explainability and reduce unauthorized cross-boundary actions, while possibly increasing approval latency and encouraging informal workarounds.

Validation should proceed in four stages. Stage 1 uses replay data and synthetic faults to test units, timestamps, missing events, drift, and mapping changes. Stage 2 uses a hardware-in-the-loop environment to test timing, fallback control, and safe recovery. Stage 3 conducts a prospective single-site pilot with pre-registered metrics and a contemporaneous baseline. Stage 4 evaluates transfer across at least two sites with different equipment and data schemas. Advancement between stages should require documented thresholds rather than qualitative demonstrations.

Logistics validation deserves a distinct set of scenarios because routing and resource-matching studies operate with assumptions that differ from plant control. Research on North American cross-border logistics, machine-learning-based network optimization, regional integration, intelligent resource matching, and standardization-led industry development provides possible scenario variables such as border delay, fragmented carrier capacity, regional platform scale, and uneven technology adoption [17][18][19][20][21]. The studies vary in empirical depth and geographical scope, and their results should not be pooled as if they measured the same outcome. They are more appropriately used to construct stress cases against which the network twin can be challenged.

Resource allocation and resilience metrics also require caution. Submodular optimization under multi-dimensional constraints offers one possible representation of spillover effects [22], while predictive safeguard architectures emphasize cascading stability across very large platforms [23]. Both perspectives may inform scenario design, but a mathematically feasible allocation can remain operationally

unacceptable, and resilience measured at platform scale may conceal losses transferred to a supplier or logistics node. Validation should consequently report node-level as well as network-level outcomes.

Dimension	Example measures	Minimum reporting requirement
Technical accuracy	State-estimation error; event-matching precision; drift-detection delay	Dataset window, ground truth, missingness, confidence intervals
Operational utility	Decision lead time; schedule stability; recovery time; constraint violations	Baseline definition, workload, disruption scenario, operator overrides
Resilience	Performance degradation and recovery under faults or delayed data	Fault injection method, fallback mode, recovery criterion
Governance	Unauthorized actions; trace completeness; approval time; policy exceptions	Decision authority, audit coverage, exception rationale

7. Discussion

The framework narrows the ambition of a "unified ecosystem" into a governed federation of purpose-bounded twins. This is a substantive change: technical integration is not treated as evidence of causal improvement, and flexibility at ingestion is not confused with semantic validity. The architecture also separates recommendations from authorized actions, which is essential when decisions cross plant, firm, and national boundaries.

Several trade-offs remain, and the same observation may admit competing interpretations. A shorter recovery time after a connectivity failure could indicate that edge autonomy is effective, but it might also reflect an unusually mild disruption, excess local inventory, or a fallback policy that postpones costs to another node. Fewer scheduling changes could signal greater stability, although they could equally reveal an optimizer that is insufficiently responsive to new information. A reduction in cross-site reconciliation errors may result from better semantics, from a narrower set of integrated variables, or from operators correcting errors outside the logged workflow. Considering these possibilities, evaluation should combine system logs, controlled fault injection, operator interviews, and sensitivity analysis rather than relying on a single performance indicator.

Potential bias may also enter through site selection and measurement design. Plants with mature instrumentation are more likely to participate in pilots and may make the architecture appear easier to transfer than it would be in legacy facilities. Operator overrides are ambiguous: they can expose model weakness, but they can also embody local knowledge that the formal model does not represent. Likewise, confidence intervals derived from historical variation may be misleading when disruptions change the underlying process. The framework remains conceptual and has not been validated in a live multi-site network; it does not establish that a particular architecture will reduce emissions, water use, logistics cost, or sensor drift. Such outcomes require domain-specific system boundaries, credible baselines, prospective evaluation, and, to some extent, institutional willingness to report unsuccessful deployments.

8. Conclusion

The deeper implication of a transnational manufacturing-network twin lies less in producing a perfectly unified representation than in making the limits of coordination visible and governable. When sensing, semantics, models, and decision authority remain distinguishable, disagreement between sites is no longer merely a data-cleaning defect; it becomes evidence about differing operational contexts, institutional priorities, and acceptable forms of risk. The proposed architecture may, to some extent, shift digital-twin research from the pursuit of representational completeness toward the design of accountable relationships among partial models.

For practice, this perspective suggests that implementation maturity should not be judged only by prediction accuracy or automation rate. The capacity to decline an unsafe recommendation, explain a semantic mismatch, preserve local operation during network failure, and document why a human overrode a model may be equally significant. Further research is needed to determine whether such safeguards remain workable under commercial pressure, cross-border data restrictions, and rapidly changing disruptions. Comparative multi-site studies, including unsuccessful or partially successful pilots, could reveal when federation genuinely strengthens resilience and when it merely redistributes uncertainty. This leads us to further thinking about digital twins not simply as technical replicas of industrial systems, but as evolving institutions through which evidence, authority, and responsibility are negotiated.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported without any funding.

Conflicts of Interest

The author(s) declare no conflicts of interest.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] *Feng, Y. (2026). Research on the Construction and Practice of the Full-Domain Architecture System for Enterprise Management Informatization under the Digital Economy. Journal of World Economy, 5(2), 17-27.*

[2] *Feng, Y. (2026). Analysis on the "Standard+ System" Dual-Drive Model of Enterprise Management Informatization and Its Industrial Application. Frontiers in Management Science, 5(2), 41-49.*

[3] *Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital Twin in manufacturing: A categorical literature review and classification. IFAC-PapersOnline, 51(11), 1016-1022.*

[4] *Kamble, S. S., Gunasekaran, A., Parekh, H., Mani, V., Belhadi, A., & Sharma, R. (2022). Digital twin for sustainable manufacturing supply chains: Current trends, future perspectives, and an implementation framework. Technological Forecasting and Social Change, 176, 121448.*

[5] *Ivanov, D., & Dolgui, A. (2021). A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. Production Planning & Control, 32(9), 775-788.*

[6] *Lee, J., Bagheri, B., & Kao, H. A. (2015). A cyber-physical systems architecture for industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18-23.*

[7] *Monostori, L. (2014). Cyber-physical production systems: Roots, expectations and R&D challenges. Procedia CIRP, 17, 9-13.*

[8] *Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637-646.*

[9] *Hao, Z. (2025). Fault-Tolerant Real-Time Scheduling for Edge AI in US Critical Infrastructure. Engineering Frontiers, 1(4).*

[10] *Hao, Z. (2026). Energy Efficient Multi Core Task Scheduling for Real Time Edge AI Systems: A Latency Aware Approach. International Journal of Advance in Applied Science Research, 5(3), 1-14.*

[11] *Balta, E. C., Pease, M., Moyne, J., Barton, K., & Tilbury, D. M. (2023). Digital twin-based cyber-attack detection framework for cyber-physical manufacturing systems. IEEE Transactions on Automation Science and Engineering, 21(2), 1695-1712.*

[12] *Meng, S. (2026). Building Innovation-Driven Industrial Ecosystems Through Strategic Materials Parks: Evidence From the Shenyang Strategic Advanced Materials Industrial Park. Development, 6, 20.*

[13] *Meng, S. (2026). Strategic Management of Large-Scale Innovation Funds: a Framework for Accelerating Technology Commercialization Across Emerging Industries. Journal of Economic Theory and Business Management, 3(2), 11-15.*

[14] *Huang, H. (2025). Development and Evaluation of a Teacher Training Program in Artificial Intelligence Technology. Journal of Advanced Research in Education, 4(1), 23-31.*

[15] *Zhang, N. (2026). Research on the Construction of a Multi-Format Integrated Management System for Small, Medium and Micro-Sized Enterprises. Frontiers in Management Science, 5(2), 9-18.*

[16] *Zhang, N. (2026). Research on the Long-Term Mechanism of Digital Transformation for Small and Medium-Sized Enterprises. Journal of Intelligence and Engineering Technology, 1(2), 19-28.*

[17] *Nie, X. (2025). Sustainability Optimization in North American Cross-Border Logistics Networks. Journal of World Economy, 4(6), 66-73.*

[18] *Shengtao, L. (2025). Machine Learning-Based Logistics Network Optimization Algorithm. Academic Journal of Computing & Information Science, 8(5), 46-54.*

[19] *Chen, Y. (2026). Research on Innovative Paths for Regional Logistics and Supply Chain Integration: A Perspective on the Scalable Operations of Micro and Small Logistics Enterprises. Journal of Intelligence and Engineering Technology, 1(1), 70-79.*

[20] *Chen, Y. (2025). Analysis of the Technical Application and Effectiveness of Intelligent Algorithms Empowering Regional Logistics Resource Matching. Engineering Frontiers, 1(3).*

[21] *Chen, Y. (2025). Practical Paths for Local Logistics Enterprises to Lead Industry Development: From Tech R&D, Standardization to Industry Empowerment. Engineering Frontiers, 1(4).*

[22] *Liu, Y. (2026). Heterogeneous Resource Slot Optimization in Multi-Dimensional Recommendation Landscapes: A Submodular Constrained Framework with Cross-Space Spillover Effects. Journal of Progress in Engineering and Physical Science, 5(1), 26-31.*

[23] *Liu, Y. (2026). Cascading Resilience Through Predictive Multi-Dimensional Safeguards: System Stability Architecture for Billion-Scale Concurrent Platforms. Innovation in Science and Technology, 5(1), 35-45.*

[24] *Meng, S. (2026). Bridging Research and Market Adoption in Artificial Intelligence: an Investment-Driven Framework for Commercializing AI Security Technologies. Academic Journal of Sociology and Management, 4(3), 12-19.*

[25] *Feng, Y. (2026). Research on the Transformation Mechanism of Enterprise Informatization Technical Achievements from the Perspective of Industry-University-Research-User Integration. Innovation in Science and Technology, 5(2), 45-53.*