The SmarTwin project, an Intelligent Digital Supply Chain Twin

Pietro Catalano[1,†], Angelo Ciaramella[2,†], Pietro D'Ambrosio[1,†], Aniello De Prisco[3,†], Michele Di Capua[4,†], Luigi Ferraro[4,5,†], Salvatore Moscariello[1,†] and Pasquale Perillo[1,†]

[1]*LinearIT spa, Via Giovanni Severano, 28, 00161 Rome, Italy.*
[2]*DIST – Dipartimento di Scienze e Tecnologie – Università di Napoli "Parthenope", Naples, 80143, Italy.*
[3]*Magsistem spa, Zona Industriale, 81030 Gricignano di Aversa (CE). Italy.*
[4]*US srl, via Porzio, Centro Direzionale di Napoli Isola G2, Naples, 80143, Italy.*
[5]*SCCC - University of Miami, 33136, Miami, USA.*

Abstract

The widespread use of artificial intelligence (AI) to extract value from historical data is transforming the supply chain landscape, adding complexity to operational management and risk mitigation. This paper provides a discussion of the SmarTwin research project and its methodology, with the aim of implementing an intelligent Digital Supply Chain Twin (iDSCT) to address the challenges of modern supply chains. The SmarTwin architecture has been designed to achieve several objectives, including establishing a traceable automated system that can improve the overall reliability of the supply chain. Furthermore, the system is intended to develop a comprehensive representation of a specific supply chain instance that serves as a controlled virtual environment to implement, simulate and optimize the decision support system (DSS). SmarTwin aggregates data and information from multiple sources and integrates them into a semantic framework. This unified vision is continuously monitored by a predictive analytics layer capable of issuing early warning alerts, which trigger coordinated responses based on combined simulation and optimization strategies to support data-driven, resilient, and sustainable supply chain operations.

Keywords

AI, supply chain, digital twin, risk management

1. Introduction

The SmarTwin project, started in 2023, is structured around a series of milestones (OR), with responsibilities distributed among research partners, and now reached its 24th month. The project has produced publications outlining its goals to investigate an innovative service model for cost optimization, risk reduction, micro-traceability of product processing steps, certification of the ecological footprint and financial support for complex supply chains [1], and inspired a focus paper on financial sustainability and the use of supply chain finance services to support it [2].

This work focuses on the Smartwin objective of developing an intelligent Digital Supply Chain Twin (iDSCT) [3] capable of replicating real-world supply chain instances by integrating diverse data sources (e.g., IoT, ERP, certified transactions based on blockchain, and financial systems) and processes from multiple stakeholders. This goal is achieved through the creation of an integration layer that enables the monitoring and analysis of critical value chain elements to mitigate disruptions. A second key objective is to ensure a high level of trustworthiness by monitoring and certifying critical operations, including workflow execution, financial transactions, and anomaly detection results. This enables transparency and traceability across all production phases, supporting asset verification from intermediate steps to final output.

Ital-IA 2025: 5th National Conference on Artificial Intelligence, organized by CINI, June 23-24, 2025, Trieste, Italy

† These authors contributed equally.

p.catalano@lineargruppo.it (P. Catalano); angelo.ciaramella@uniparthenope.it (A. Ciaramella); p.dambrosio@lineargruppo.it (P. D'Ambrosio); nello.deprisco@magsistem.com (A. D. Prisco); m.dicapua@usmail.it (M. D. Capua); s.moscariello@lineargruppo.it (S. Moscariello); p.perillo@lineargruppo.it (P. Perillo)
0009-0001-0644-4388 (P. Catalano); 0000-0001-5592-7995 (A. Ciaramella); 0000-0001-8904-180X (M. D. Capua)

To achieve these goals, the system was designed around pivotal non-functional requirements. It adopts a decentralized, Web3-based, event-driven architecture, and comprises loosely coupled, specialized subsystems for scalable and flexible configurations. The subsystems are responsible for distinct functional domains, and can be added or removed as needed. The macro-functionalities of SmarTwin emerge from the coherent integration of the capabilities of the active subsystems within the current configuration.

The AI Subsystem represents a crucial component of the system, offering a pluggable set of AI tools delivered as a service that support core functionalities to address various supply chain challenges.

Section 2 presents the architectural design and its core elements. Section 3 presents the AI as a service approach and its role in enabling system functionalities.

2. Architectural design

This section outlines the architectural design principles adopted to support the system operations and meet key non-functional requirements. As shown in Figure 1, the architecture comprises infrastructural components (UML components) and specialized subsystems (UML packages), integrated in a decentralized manner through the Blockchain Subsystem, which supports the event-driven supply chain workflow. The following high-level overview describes the main elements of the architecture, with reference numbers corresponding to those in the figure (noted in parentheses).

The *Generic Stakeholder* component (1) is intended to abstract any participant in the supply chain who contributes to the realization of the final product and shares data related to specific transformation phases. It provides a multitude of data types, including silos data (ERP, financial), operational data, formal declarations, IoT, and video streams.

The *Data Lake* component (2) stores data shared by stakeholders and any information produced by the SmarTwin subsystems. The employment of AI functionalities facilitates the classification of data [4] according to provided taxonomies. The technological solution is based on the Hadoop framework[1].

The *IoT-queue* component (3) is a middleware that collects the data provided by IoT devices installed in the work areas of stakeholders. These measurements are then integrated into the digital twin of the corresponding supply chain instance. It supports various communication protocols for non-invasive interoperability with stakeholder infrastructures, thereby ensuring seamless integration. It is implemented using Eclipse Hono[2].

The IoT Manager Subsystem (4) is responsible for processing IoT-Queue sensor data to build a semantic representation that enables the creation of a digital twin for a supply chain instance. It is also responsible to manage real-time alerts and prepare time-series data for machine learning tasks. It uses Eclipse Ditto[3] and InfluxDB[4].

The *Industrial Knowledge Management* component (5) ingests data from the Data Lake component to construct a financial knowledge graph, assessing supply chain and stakeholders financial health and risk. The use of AI techniques facilitates the extraction and inference of knowledge from data [5]. The implementation of the component is based on Neo4j[5], a graph database management system.

The *Visual Data Manager Subsystem* (6) is responsible for retrieving the video streams shared by the stakeholders. It is also tasked with the preparation and analysis of them using computer vision models that are provided by the AI Subsystem. The results are transmitted to the IoT queue to enrich the digital twin managed by the IoT Subsystem.

The *Blockchain Subsystem* (7) is designed to handle data certification and the decentralized logging of system and workflow events, supporting the system's Web3 and event-driven foundation [6]. The

[1]https://hadoop.apache.org
[2]https://eclipse.dev/hono
[3]https://eclipse.dev/ditto
[4]https://www.influxdata.com
[5]https://neo4j.com

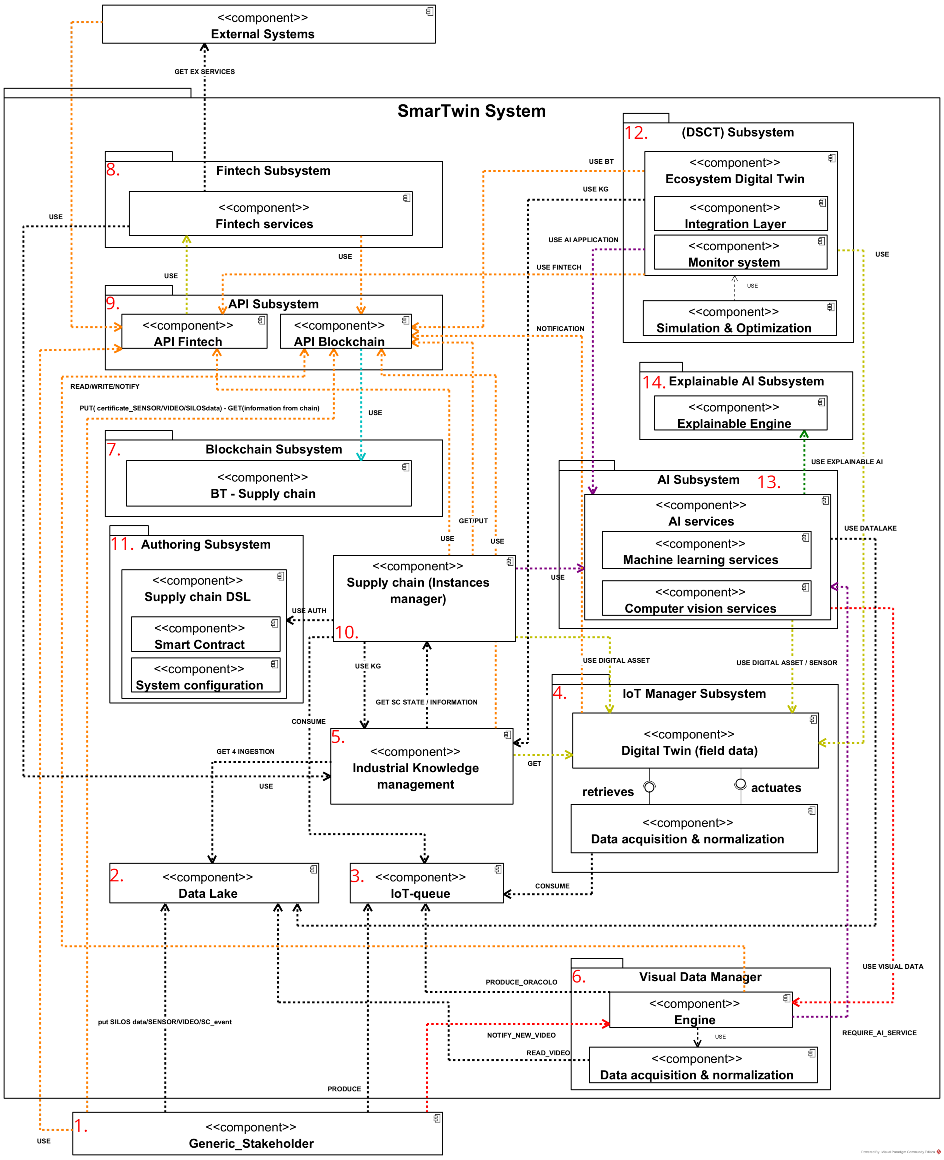

Figure 1: SmarTwin project architecture.

distributed network is implemented using Hyperledger Besu[6], a permissioned blockchain, and Solidity smart contracts.

The *Fintech Subsystem* (8) is a portal that offers financial technology services based on credits derived from supply chain transactions [7]. For example, stakeholders can securitize matured credits and convert them into tradeable assets to obtain immediate liquidity and reduce financial disruptions.It relies on blockchain subsystem certified data and and is accessible to external entities like credit institutions (shown as External Systems in the component diagram).

The *API Subsystem* (9) offers standardized, secure REST interfaces to access blockchain and fintech services. A technological solution is currently under development, based on Hyperledger FireFly[7] and aligned with Open Banking guidelines.

The *Supply Chain instance manager* (10) is responsible for managing workflow execution, data verification, and product quality assessment, managing multiple SC instances concurrently via blockchain event monitoring.

The *Authoring Subsystem* (11) is responsible for building the system configuration. Provides tools for

[6]https://besu.hyperledger.org
[7]https://hyperledger.github.io/firefly

domain experts to manage the complexity of the task minimizing the need for developer support. A domain-specific language (DSL) is under development to configure supply chain instances and define related smart contracts in Solidity for the Blockchain Subsystem. The use of a fine-tuned LLM is in consideration to translate contract requirements into the Solidity DSL through a conversational interface.

The *DSCT Subsystem* (12) provides a holistic view of the entire supply chain by aggregating data from all components and subsystems. It establishes a semantic integration layer that is used to generate multiple views highlighting key aspects such as financial performance, product quality, carbon footprint, and ethical considerations, among others. Its core function is to enable proactive monitoring through an AI-powered prediction layer that supports real-time analysis, identifies critical data paths, and forecasts potential scenarios, such as financial disruptions. Upon detecting risks, the subsystem issues early warnings to the monitoring system, triggering contingency management through simulation and optimization. These processes yield prescriptive suggestions to support decision-making. The subsystem also includes a suite of Grafana-based[8] dashboards, providing a comprehensive visualization of interconnected components and the current state of supply chain.

The *AI Subsystem* (13) manages and delivers as a service AI models (machine learning, computer vision) to other subsystems and components of SmarTwin, exposing them as REST/gRPC microservices using the Seldon Core framework[9]. provides more information on its AI as a service approach.

The *Explainable AI Subsystem* (14) provides tools to interpret black-box AI decisions and translate them into human-understandable explanations, improving transparency and stakeholder trust.

3. The "AI as a service" approach in Smartwin

This section discusses the application of artificial intelligence and data analysis techniques in the development of SmarTwin functionalities. Table 1 summarizes the methodologies that will be implemented, organized by subsystem and infrastructural component. These techniques support key system operations, as detailed below.

Table 1
AI/data analysis techniques in SmarTwin functionalities

AI/Data Analysis Methodologies	Used by	Description
Computer Vision models (Vision Transformer, YOLO)	Visual Data Manager	Provide quality inspection functionalities and condition assessment of assets.
Machine learning models (time series analysis, anomaly detection)	IoT Manager Subsystem	Detect patterns in real-time digital twin data to predict problematic scenarios (e.g., predictive maintenance).
Knowledge Graph Reasoning (relational inference via embeddings and graph-based learning)	Industrial Knowledge Management	Predict potential relationships in the knowledge graph ingestion pipeline.
Graph Data Science (Graph-based inference and algorithms)	Industrial Knowledge Management	Analyze graph structure (e.g., node roles, communities) to extract insights.
Text Classification and Labeling (Supervised Learning)	Data Lake	Data classification and semantic enrichment using predefined taxonomies.
Natural Language Processing (LLMs, Text Generation)	Authoring Subsystem	Generation of DSL description for supply chain instances configuration.
Machine Learning models (Predictive analytics and forecasting)	DSCT	Implement the prediction layer that analyze DSCT integrated data to anticipate and manage potential risks.

A main project goal is the construction of a prediction layer for real-time monitoring and risk

[8]https://grafana.com
[9]https://www.seldon.io/solutions/core

identification within the DSCT subsystem to provide proactive suggestions about anomalies and potential disruptions [8]. This layer will use a set of AI models and operate on an integrated data layer built from the multiple subsets of system data. At the current stage of the project, the required datasets have not yet been collected, so no final technology choices for these models have been made and experiments will be conducted before the project's conclusion. However, the AI Subsystem is designed to remain flexible, hosting a repository of models that can be updated with minimal integration effort. This allows for a highly customizable iDSCT, adaptable to the specific supply chain instance being monitored.

The concept of AI as a Service (AIaaS) refers to the delivery of artificial intelligence capabilities as modular, API-accessible services, hosted on cloud or hybrid infrastructure, and consumable on demand by external applications. This paradigm enables organizations to integrate intelligent capabilities, such as classification, anomaly detection, prediction, or language processing, without having to deal directly with managing the entire machine learning pipeline, from development to deployment.

For SmarTwin, which operates in complex scenarios like supply chain control in the transport of perishable fruits and vegetables, adopting a Seldon Core-based deployment architecture offers key advantages. These include timely, scalable, and reliable model orchestration—especially for use cases involving computer vision, time series analysis, and predictive maintenance.

At this late stage of the project, it has been possible to appreciate one of the main advantages of using Seldon Core: its ability to orchestrate ML models in Kubernetes environments, providing a flexible, cloud-native infrastructure that facilitates the deployment and management of heterogeneous models in complex manufacturing environments. In contexts such as agribusiness, where data can come from multiple sources, such as images acquired by cameras for automated visual inspection of product surfaces and IoT sensors for supply chain monitoring [9], Seldon's modularity enables the integration of distributed inference pipelines, supporting pre- and post-processing models, custom transformations, and ensemble logic. In mission-critical applications such as the early prediction of failures in container refrigeration systems or the identification of anomalies during transportation phases, the ability to detect conceptual changes or performance degradation in real time enables timely intervention and improves the resilience of the supply chain.

Seldon Core also offers robust model versioning, testing, release and rollback mechanisms for services, which are key aspects in scenarios where models need to be frequently updated or new ones added without compromising service continuity. This is particularly relevant when considering models that are adaptive or subject to periodic retraining based on recent data, as in the case of visual quality classification of vegetables and fruits, which is subject to seasonal and environmental variations.

However, despite the benefits, its adoption also presents challenges. Integration demands advanced Kubernetes and DevOps knowledge, increasing initial project complexity and slow down prototyping phases, especially when rapid deployment is required. Additionally, managing distributed components and maintaining extensive YAML configurations introduces technical overhead that is unnecessary in simpler or low-scale scenarios. Another limitation involves latency, introduced by containerization and the orchestration of complex pipelines. In applications requiring real-time inferences with sub-second response times, as quality checks on high-speed packaging lines, the architectural overhead may necessitate a trade-off between model accuracy and speed of response.

Currently, we believe that Seldon Core is a powerful and scalable solution for deploying artificial intelligence models in highly complex industrial settings such as the transportation of perishable fruit and vegetables. However, its success depends on a proper assessment of the trade-off between architectural power and operational complexity, as well as a harmonious integration with the digital ecosystem of the agrifood supply chain.

4. Conclusions

SmarTwin is a research project aimed at addressing complex challenges in the supply chain domain through the development of a configurable and scalable digital solution. At its core lies an intelligent digital twin of the supply chain, designed to ensure certification, traceability, and transparency across all operational stages. The system integrates a holistic view of supply chain dynamics with a predictive analytics layer capable of monitoring anomalies and forecasting potential disruptions.

Following 24 months of research and development, the project has now entered its final phase, during which validation, integration, prediction and optimization activities will be completed.

Acknowledgements

The project was funded by the Italian Ministry of Economic Development (MISE): it started in 2023 and it will end in 2026. Project reference no. F/310218/05/X56.

Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-4o and DeepL in order to: grammar and spelling check, paraphrase, reword and improve writing style. The author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] M. D. Capua, E. D. Nardo, A. Ciaramella, G. Iannuzzo, A. D. Prisco, D. R. Laderchi, P. Catalano, P. D'Ambrosio, S. Moscariello, Design the modern supply chain: The smartwin project, in: S. D. Martino, C. Sansone, E. Masciari, S. Rossi, M. Gravina (Eds.), Proceedings of the Ital-IA Intelligenza Artificiale - Thematic Workshops co-located with the 4th CINI National Lab AIIS Conference on Artificial Intelligence (Ital-IA 2024), Naples, Italy, May 29-30, 2024, volume 3762 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024, pp. 212–217. URL: https://ceur-ws.org/Vol-3762/479.pdf.

[2] G. Galante, C. Esposito, P. Catalano, S. Moscariello, P. Perillo, P. D'Ambrosio, A. Ciaramella, M. D. Capua, Financial twin chain, a platform to support financial sustainability in supply chains, 2025. URL: https://arxiv.org/abs/2503.15980. arXiv:2503.15980.

[3] D. Ivanov, Intelligent digital twin (idt) for supply chain stress-testing, resilience, and viability, International Journal of Production Economics 263 (2023) 108938. doi:10.1016/j.ijpe.2023.108938.

[4] Y. He, J. Song, Y. Wang, S. Chaudhuri, V. Anil, B. Lassiter, Y. Goland, G. Malhotra, Auto-tag: Tagging-data-by-example in data lakes, CoRR abs/2112.06049 (2021). URL: https://arxiv.org/abs/2112.06049. arXiv:2112.06049.

[5] H. Zhong, D. Yang, S. Shi, L. Wei, Y. Wang, From data to insights: the application and challenges of knowledge graphs in intelligent audit, Journal of Cloud Computing 13 (2024). URL: http://dx.doi.org/10.1186/s13677-024-00674-0. doi:10.1186/s13677-024-00674-0.

[6] M. Rabe, H. Rainer Körsgen, Multi-layer track & trace system landscape with smart contract validation instances, Journal of Data Science and Intelligent Systems (2024). URL: https://ojs.bonviewpress.com/index.php/jdsis/article/view/3384. doi:10.47852/bonviewJDSIS42023384.

[7] S. Chakuu, D. Masi, J. Godsell, Exploring the relationship between mechanisms, actors and instruments in supply chain finance: A systematic literature review, International Journal of Production Economics 216 (2019) 35–53. URL: https://www.sciencedirect.com/science/article/pii/S0925527319301367. doi:https://doi.org/10.1016/j.ijpe.2019.04.013.

[8] A. Aljohani, Predictive analytics and machine learning for real-time supply chain risk mitigation and agility, Sustainability 15 (2023). URL: https://www.mdpi.com/2071-1050/15/20/15088. doi:10.3390/su152015088.

[9] S. Deng, H. Zhao, W. Fang, J. Yin, S. Dustdar, A. Y. Zomaya, Edge intelligence: The confluence of edge computing and artificial intelligence, IEEE Internet of Things Journal 7 (2020) 7457–7469. doi:10.1109/JIOT.2020.2984887.