inetum.

WHITE PAPER

Convergence of AI and Digital Twins in telecom

A closed loop: How AI automates simulations and generates synthetic data to drive predictive models

A closed loop: How AI automates simulations and generates synthetic data to drive predictive models

Table of Contents

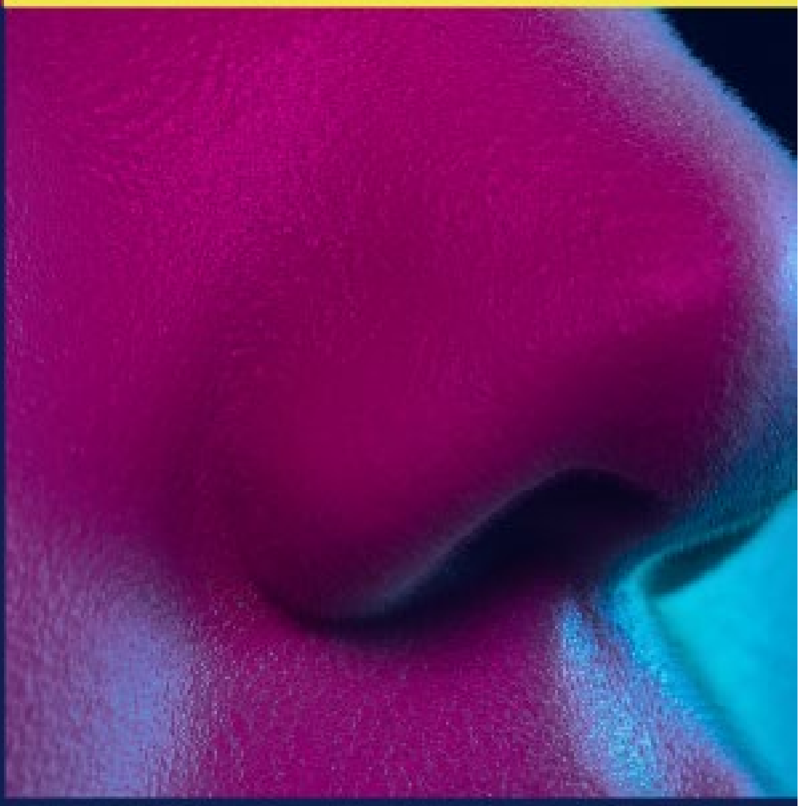

01

Introduction

1 Introduction

A Digital Twin is a virtual representation of a real-world object, system, or process

A **Digital Twin** is a virtual representation of a real-world object, system, or process. In telecommunications, it means creating a digital copy of a network, service, or infrastructure that can simulate behavior, anticipate issues, and support better decision-making. It's like having a virtual lab where different scenarios can be tested without impacting the live environment, accelerating innovation and improving operational efficiency.

This document combines the TM Forum's **DT4DI framework** (Digital Twin for Decision Intelligence) with advanced AI techniques to build a **closed loop** for simulation and machine learning. The initial phase adapts the DT4DI v2.2.0 **reference architecture**, integrating an automated simulation "engine" (according to TR284G) and a data **pipeline** that feeds predictive AI models (IG1310A). The result is a platform capable of launching thousands of simulations, validating results, and continuously updating AI models, thereby improving the accuracy of predictions (e.g., for energy consumption or anomaly detection). Inetum, aligned with the standards and reference architectures of TM Forum (DT4DI v2.2.0, ODA, AN—Autonomous Networks), recommends this base platform as the cornerstone of any Digital Twin + AI project in Telecommunications.

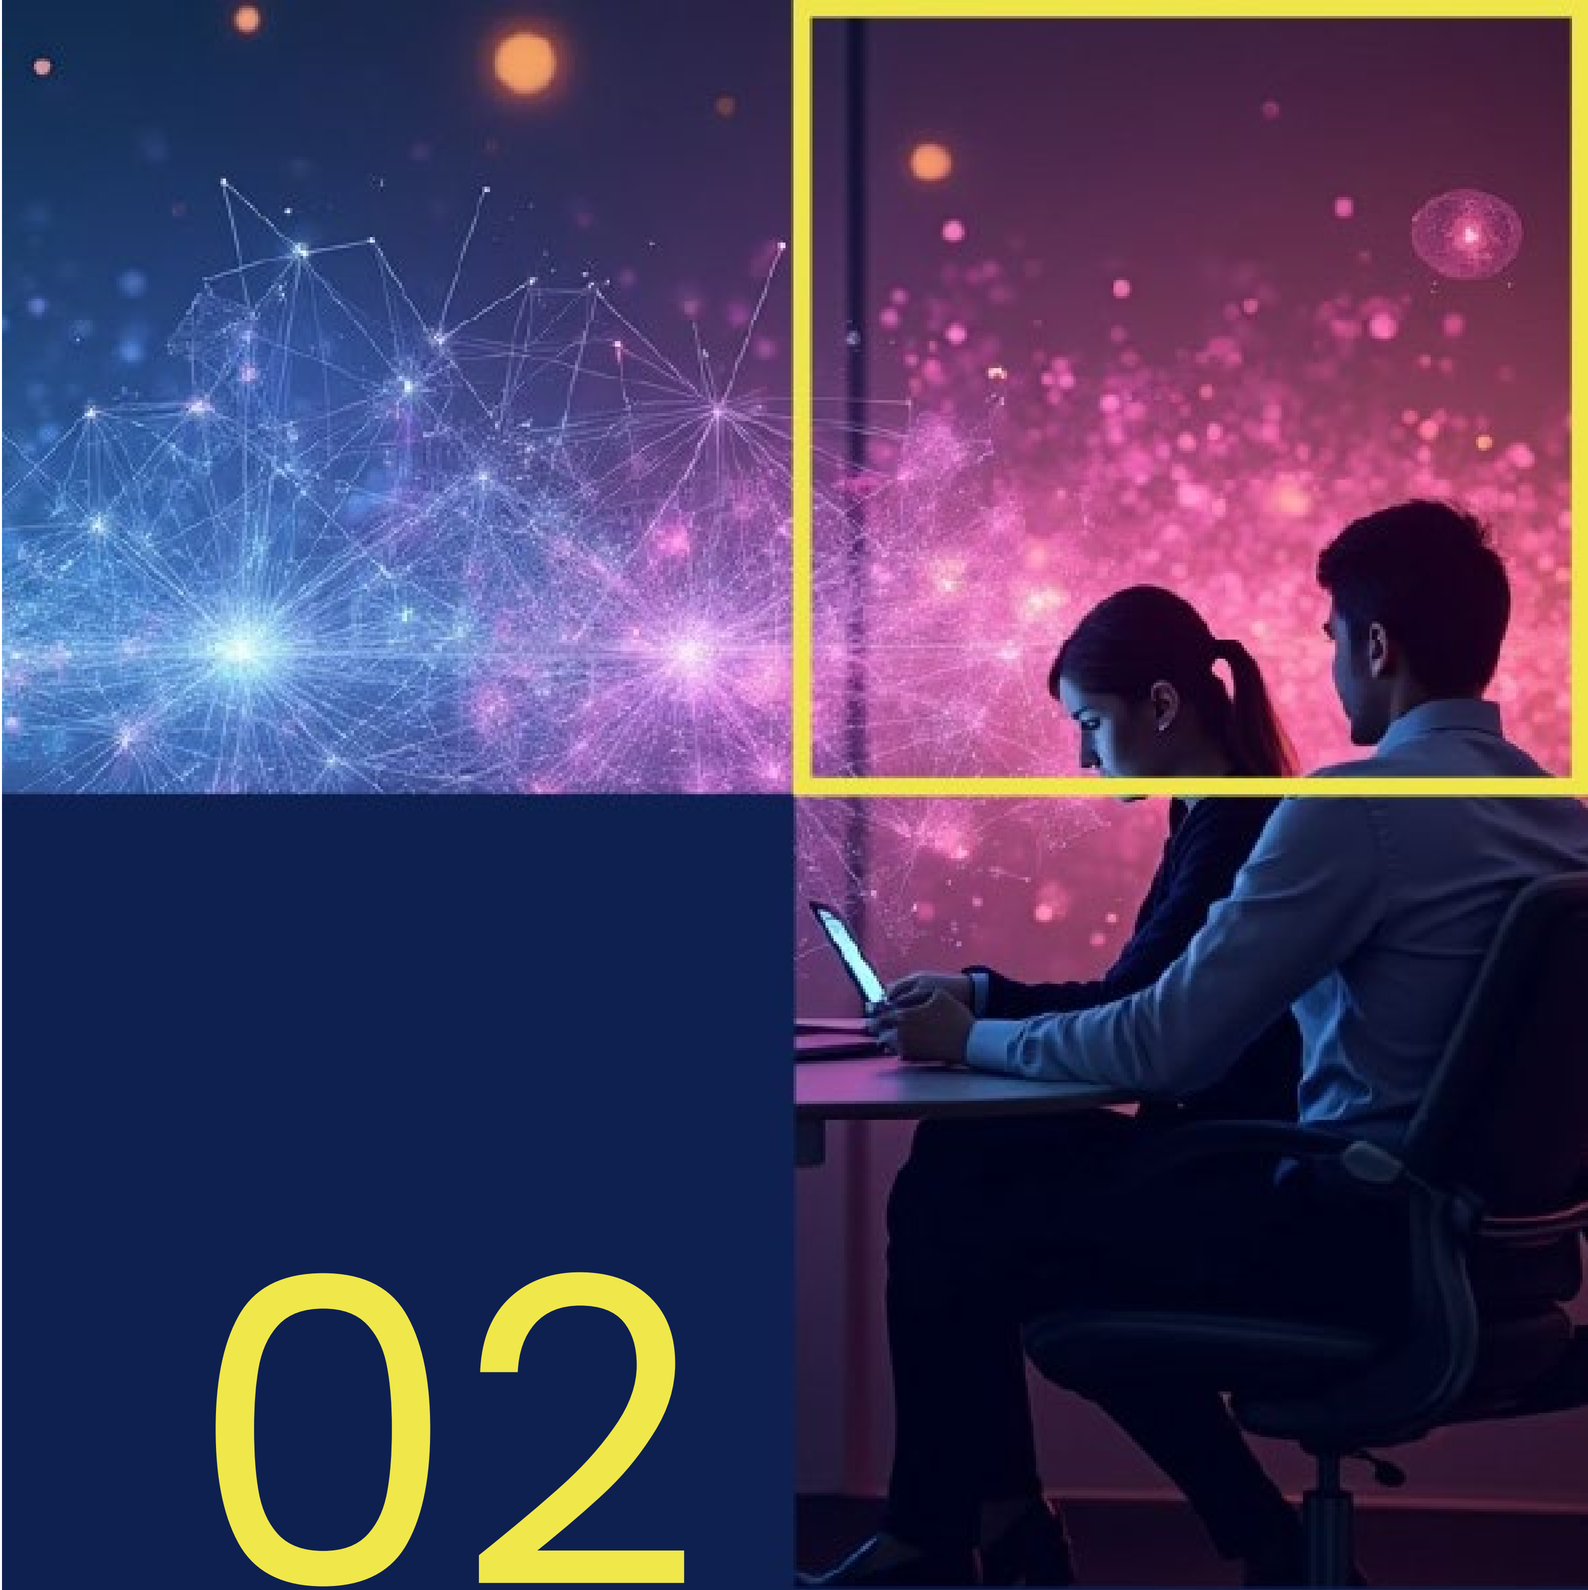

02

Automation of simulations without human intervention

2 Automation of simulations without human intervention

- In the vertical process Methodologies view, as illustrated in Figure 1, **DTOps** (Digital Twin Operations) are introduced alongside **AIOps** and **DIOps**. These are methodologies that automatically manage the lifecycle of digital twins, AI components, and decision processes, respectively, enabling a closed loop of continuous simulation, validation, and adjustment without manual intervention.

- In the **DT Modeling layer**, models incorporate "**simulation feedback**" as a data source to refine the virtual replicas, integrating real-time data, AI models, and simulation results to programmatically generate new runs.

- Within **DI Enabling Services**, one of the essential services is simulation. It is classified alongside descriptive, diagnostic, predictive, and prescriptive services, which implies that these simulations can be invoked via API and executed in a fully automated manner.

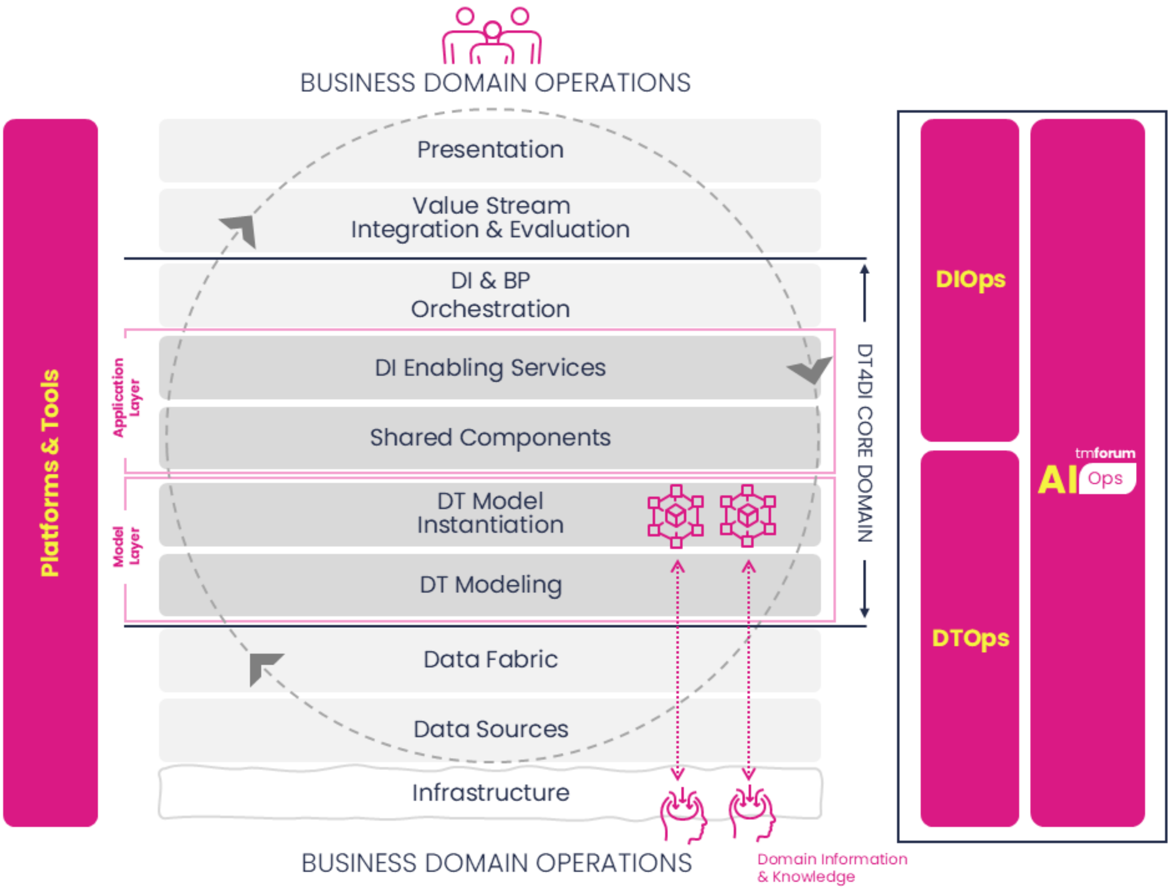

Figure 1: DT4DI Framework, including the Level 0 DT4DI reference architecture, layer view

03

Data and Behavior Feedback for Training Predictive Models

3 Data and Behavior Feedback for Training Predictive Models

- The **Data Fabric** and the **DT Modeling layer** configure pipelines that capture both real operational data and simulation results, normalizing and storing them in repositories accessible for AI training.
- The **AIOps** cycle is designed to manage not only the deployment but also the **continuous training** and updating of AI models based on new data and behaviors extracted from the digital twins, thus improving the accuracy of predictions in production.
- In the **BP & DI orchestration**, an explicit **"feedback loops"** step appears, where the results of decisions and simulations feed back into both the AI models and the business processes, closing the machine learning circuit.

04

Frameworks and Maturity for Implementing AI + Digital Twins

4 Frameworks and Maturity for Implementing AI + Digital Twins

DT4DI Reference Architecture stands as the **Domain Architecture** that unifies DT, AI/ML, and Decision Intelligence, and is complemented by three vertical **frameworks**:

AIOps Toolkit (AI management) is the set of practices, processes, and principles recommended by TM Forum to govern the entire lifecycle of AI components within a DT4DI solution. Its objectives are:

- Design and development:
  - Define quality, robustness, and explainability standards for AI models.
  - Incorporate automated regression and training data validation tests.
- Deployment and execution:
  - Orchestrate CI/CD pipelines for models (training, validation, promotion to production).
  - Monitor real-time performance (latency, throughput) and business metrics (precision, recall).
- Maintenance and continuous improvement:
  - Orchestrate Detect data drift or model degradation through automated alerts.
  - Schedule model retraining and recalibrations based on quality thresholds.

DTOps (digital twin management) is the methodology that covers the operation and governance of digital twins (DT) in production. Its main blocks are:

- Definition and deployment:
  - Configuration of pipelines that feed the DT with real (real/near-real time) data and with automatically generated "simulation feedback".
  - **Model Deployment** strategies: sizing, autoscaling, and fault tolerance for DTs.
- Synchronization and consistency:
  - Twin fidelity metrics (fidelity score) comparing virtual vs. real behavior.
  - Automatic processes for model reconciliation and refinement to improve their accuracy with each cycle of simulation and real data.

- Monitoring and maintenance:
    - Versioning of logical specifications (physical models, ontologies) and their mapping to container or virtual machine infrastructures.
    - Use of specialized databases (graph DB, time-series DB, NoSQL) to ensure low latency and consistency in the replicas.

DIOps (decision process management) orchestrates the implementation and lifecycle of automated decision processes:

- Decision flow design:
    - Formal definition of decision points, criteria, inputs, and outputs, using DMN or CDD notations.
    - Catalog of DI services (descriptive, diagnostic, predictive, prescriptive, and simulation) invocable via API.
- Orchestration and integration:
    - Modeling of BP & DI workflows, where critical points trigger DI services and then reintegrate the decision into the business process.
    - Creation of feedback loops to feed the results of each decision back into both the AI models and the business rules.
- Governance and continuous improvement:
    - Monitoring of decision KPIs (time to decide, adoption rate, recommendation quality).
    - Dynamic adjustment of thresholds and criteria based on effectiveness and risk analysis, closing the machine learning loop.

These frameworks define the design, deployment, monitoring, and governance guidelines necessary to bring to production and scale digital twin solutions with AI. Furthermore, it aligns with other TM Forum standards (IG1310 Ontology, IG1307 Whitepaper) to assess the organization's **maturity** regarding the integration of digital twins and AI, and to draw an evolution roadmap.

05

Base Platform for Massive Network Data Collection

5 Base Platform for Massive Network Data Collection

To enable large-scale digital twins and AI/ML applications in the telecommunications field, it is essential to have a five-layer data architecture designed to ensure flexibility, performance, and governance. The following base platform architecture is the **recommendation of Inetum Consulting**, based on TM Forum standards.

5.1 Infrastructure

The infrastructure constitutes the physical and logical support on which data services are built. It must allow for the scalable execution of pipelines, the management of large volumes of data, and real-time operation. To do this, it requires:

- **Distributed compute,** through orchestrated containers (e.g., Kubernetes), that supports variable and simultaneous workloads.
- **Scalable storage**, combining distributed file systems, NoSQL databases, and object stores.
- **Low-latency network**, essential for absorbing network events, IoT signals, and 5G traffic without degrading the experience. Cases like that of BT Group, which has migrated its systems to Google Cloud and implemented modular components on the Open Digital Architecture (ODA), demonstrate that an "**AI-ready**" infrastructure is key to supporting network intelligence models.

5.2 Data Sources

The data sources layer integrates information from multiple systems and domains. It includes structured and unstructured data from legacy systems (BSS/OSS), network telemetry, IoT sensors, digital channels, as well as external data from third parties or open platforms. The growing complexity and heterogeneity **demand decentralized and domain-oriented** architectures, where each area (such as RAN, core, sales, or customer service) manages its flows as "data products." This model promotes accessibility, improves data quality, and enables localized knowledge that enhances advanced analytics.

5.3 Data Fabric

The data fabric acts as a transversal layer that virtualizes, orchestrates, and governs all data flows to and from the platform. Its mission is to ensure unified, consistent, and secure access to all datasets, regardless of their origin. The key elements of this layer include:

- **Intelligent metadata catalogs**, which allow for tracking data quality and lineage.
- **Ontologies and knowledge graphs**, which provide semantic meaning and context to data in complex domains like telecommunications.
- **Centralized governance and security policies**, which guarantee privacy and responsible use, especially in sensitive datasets such as customer data. The use of FAIR principles (Findable, Accessible, Interoperable, Reusable) and the ability to orchestrate real-time data enable both automated decision-making processes and the operation of digital twins.

5.4 Digital Twin Modeling (DT Modeling)

Digital twins require a logical modeling layer that faithfully represents the physical, virtual, and service elements of the network. Based on the TM Forum's DT4DI framework, it is recommended to:

- Use **graph-based models** to represent dynamic relationships between nodes, links, services, and customers.
- Apply **federated ontologies** that allow for a shared semantics between the different technical and business domains.

  These models are integrated with AIOps layers to enable complex simulations, prediction and optimization use cases, as well as automated decision engines in cognitive networks.

5.5 Digital Twin Instantiation (DT Model Instantiation)

Once defined, the models must be deployed as active virtual replicas that remain synchronized with the real systems. This layer implies:

- **Flexible deployment** using containers or virtual machines, with automatic scaling according to simulation or analysis demand.
- **Specialized storage engines** such as **graph DBs** (for topological relationships), time-series DBs (for metrics), and NoSQL (for configurations and states).
- **Unified programmatic interfaces (APIs)** that abstract the technical complexity of the back-end and allow integration with AI engines or operational dashboards. Cases like that of BT Group illustrate the challenge of synchronizing real-time data from old and modern systems, and standardizing the data pipelines necessary to keep the digital twins updated and reliable.

06

Digital Twin in 5G Networks

6 Digital Twin in 5G Networks

A very powerful use case would be to build a digital twin of the 5G network to generate **synthetic operational data** (e.g., coverage maps, quality of service metrics, mobility patterns) to feed and train AI models for network optimization. Since **5G is a very recent technology**, large volumes of historical data are not yet available in many deployments; therefore, the digital twin serves as an **additional source** of representative data, ensuring sufficient variety and volume to automate and refine predictive AI models.

6.1 Data Sources Synthetic Data Generation

Once the twin is built with the available information (city maps, antenna locations, initial site surveys), the system can automatically sweep through thousands of different scenarios (different antenna configurations, traffic variations, weather conditions), producing massive volumes of performance data that would not exist on the real network.

This synthetic data covers "dead zones" or rare extreme conditions and enriches the available historical data, becoming a **data seedbed** for training AI algorithms for **self-optimizing networks** (SON), congestion prediction, or predictive maintenance:

- **Bootstrap with physical and static models:** initially, the twin is fed with physical constraints (propagation models, 3D geodata, antenna parameters) and a small amount of **real** measurement **data** (drive-tests, KPIs).
- **Progressive learning:** as the twin generates synthetic data and it is compared with new real measurements, the models are refined (the "**simulation feedback**" loop) to improve the fidelity of the virtual replica and the quality of the generated data.
- In this way, even though 5G has little historical data, the twin acts as both a data source and validator, continuously feeding both the predictive AI model and the evolution of the twin itself.

6.2 Data Key Advantages

- **Data scalability:** thousands of hours of simulation in minutes.
- **Total coverage:** scenarios that would be difficult to measure live.
- **Controlled environment:** "what-if" experimentation before touching the real network.
- **Greater AI robustness:** models trained with more diverse and complete data.

07

Real Cases

7 Real Cases

7.1 Data Key Advantages Ericsson + NVIDIA Omniverse

In the real use case of Ericsson with NVIDIA Omniverse and O-RAN architectures, the synergy of GPU-accelerated simulation, generative AI, and synthetic data generation has made it possible to automate 5G network validation and optimization processes at a city scale, as well as to train and validate predictive models in a virtual "sandbox" before their deployment in production.

This project arose from the **increasing complexity of 5G deployment** and the challenges presented by the RAN (Radio Access Network). With the proliferation of microcells and connected devices, the interactions of radio waves with the urban environment (buildings, vegetation, vehicles) become critical to ensuring optimal coverage and quality of service. Traditional methods, based on on-site tests or overly simplified simulators, lengthen deployment cycles and do not reflect real conditions with sufficient precision, leaving a margin of error in dynamic urban environments.

To overcome these challenges, the **solution's architecture** combined two key components:

- NVIDIA Omniverse Platform with RTX GPUs:
    - Ericsson integrated its proprietary 5G propagation simulation engine with NVIDIA Omniverse Enterprise. This allowed them to leverage real-time ray tracing to model signal-environment interactions at a city scale with high fidelity, drastically accelerating design validations and the review of complex scenarios.
- Ericsson Site Digital Twin (ESDT) + Generative AI:
    - **Massive telemetry:** ESDT receives and processes real-time data from hundreds of base stations during peak demand, capturing performance metrics, energy consumption, and configuration parameters
    - **Generative AI:** machine **learning** algorithms analyze traffic patterns and user behavior to predict congestion and recommend automatic adjustments to parameters such as antenna tilt, **beam** alignment, and radio resource allocation.
    - **KPI Automation:** in collaboration with Ericsson Network Intelligence (ENI), delay and error rate thresholds are automatically updated to maintain the quality of service within defined parameters, without the need for manual intervention.

From these components, **mechanisms for generating and using synthetic data** were built to enrich the real datasets:

- **"What-If" Scenarios in O-RAN:** The digital twin creates thousands of configuration variants, such as different antenna positions, changing environmental conditions, and variable traffic patterns, and generates synthetic datasets that cover rare or extreme use cases not frequently observed in real network data.

- **Training of predictive models:**
  - **Energy optimization:** machine learning models trained with synthetic data generate predictions of base station energy consumption and define power on/off policies for RRUs (Remote Radio Units) according to hourly and geographical demand.
  - **Predictive maintenance:** simulations of component degradation (fans, power supplies, RF elements) generate failure data that improve the precision of early anomaly detection algorithms, anticipating repairs before an actual failure occurs.

This entire infrastructure has led to several tangible **benefits and results**:

- **Increased capacity:** During traffic peaks, a capacity increase of up to **14%** was achieved, without direct manual intervention on the network.
- **Cost and risk reduction:** By validating configurations and changes in a virtual **sandbox**, production update risks were reduced, lowering costs associated with **post-deployment** corrections.
- **Time-to-market agility:** Thanks to the acceleration by RTX GPUs and ray tracing in **Omniverse**, dozens of design and optimization variants were tested in minutes or hours, compared to the days or weeks that traditional methods consumed.

This approach of **urban digital twin + generative AI + synthetic data** has positioned Ericsson and its operator clients at the forefront of 5G network innovation, demonstrating how O-RAN architectures can benefit from advanced virtual environments to reduce time, costs, and improve quality of service in complex urban environments.

7.2 Use Case: Huawei – Intelligent predictive maintenance in 5G networks using digital twins and AI

Huawei has developed a predictive maintenance solution based on digital twins, combining IoT sensors, advanced analytics, and artificial intelligence to optimize the operation of critical infrastructures in networks and data centers. This solution has been applied, among other environments, in the Moro Hub data center (United Arab Emirates), managed through Huawei's modular platform.

By creating digital replicas of physical components, such as power systems, cooling, servers, and network elements, and instrumenting them with real-time sensors, the platform allows monitoring the health status of each asset and anticipating possible anomalies. The collected data (temperature, vibration, load, humidity, etc.) are continuously analyzed in the cloud or on edge environments using AI models trained to identify failure patterns.

This architecture allows for:

- Predicting critical events before they occur and making scheduled maintenance decisions.
- Optimizing the use of technical resources by reducing unnecessary interventions.
- Visualizing the complete state of the system through the centralized digital twin

Additionally, Huawei has explored collaboration with companies like **GE Digital** to develop industrial solutions that integrate predictive maintenance, digital twins, and artificial intelligence in industrial and telecommunications environments.

Although specific quantitative improvement metrics have not been published, the approach shows a clear commitment to the intelligent automation of operation and maintenance as a way to improve reliability, reduce operational risks, and advance towards **self-managing** networks.


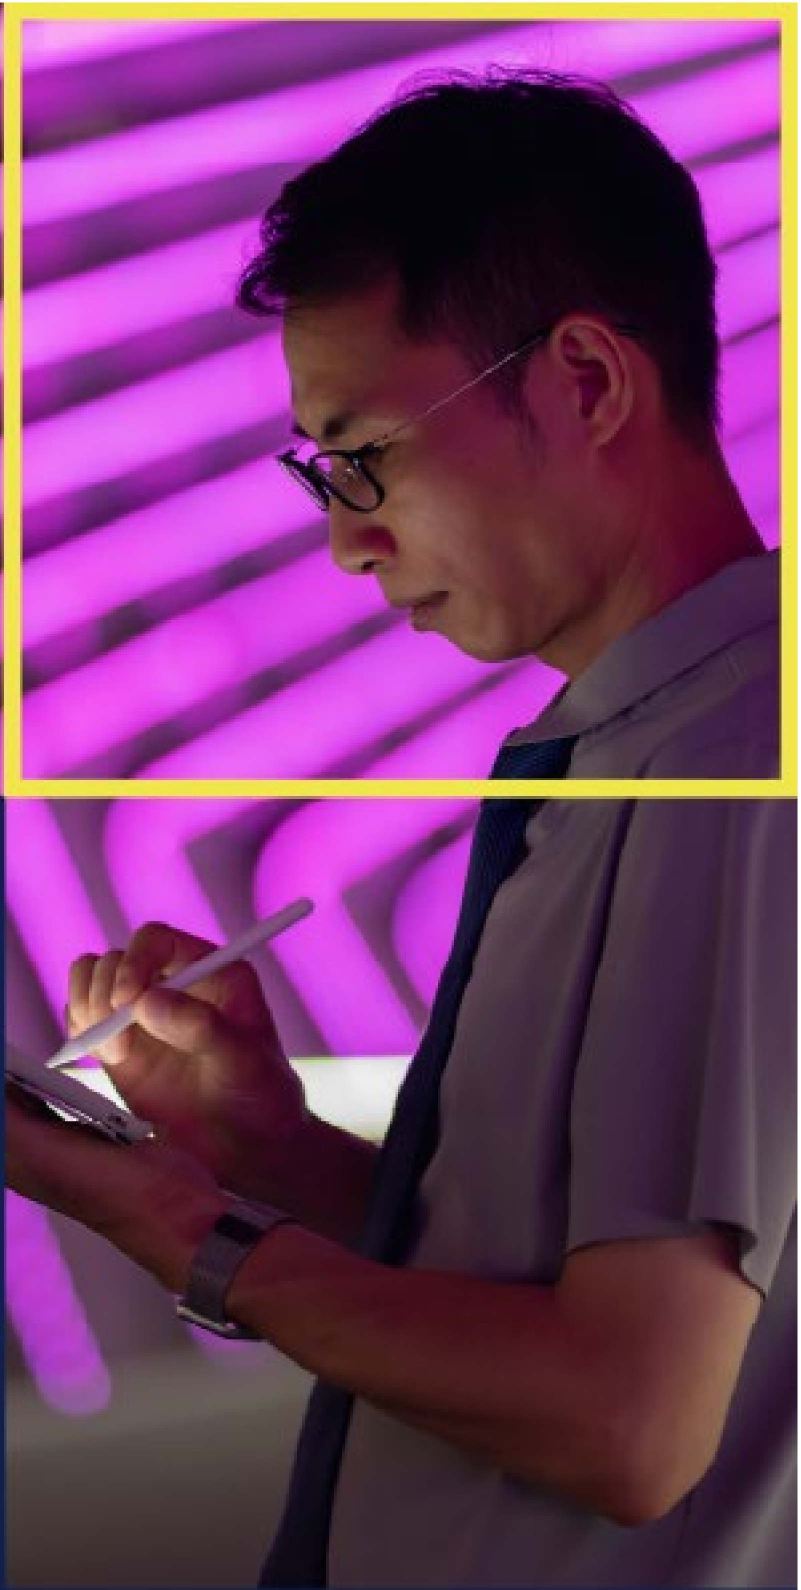

08

Conclusion and next steps

8 Real Cases Conclusion and next steps

The integration of artificial intelligence into digital twins is redefining how telecommunications networks are designed, optimized, and operated. This approach provides a safe and dynamic virtual environment where it is possible to validate strategies, anticipate failures, and train predictive models without directly impacting production operations.

Beyond accelerating decision-making, tangible benefits are beginning to emerge in key areas such as resilience, operational efficiency, and sustainability. For example, organizations that have implemented predictive maintenance enabled by digital twins report:

A **30–50%** reduction in unplanned downtime.

A **20–25%** increase in asset lifespan.

A **10–20%** decrease in maintenance costs.

A **25%** reduction in latency compared to static methods.

These results reflect the real and quantifiable potential of the Digital Twin + AI approach. In a context where the complexity of 5G networks continues to grow, companies that adopt these models will be better positioned to lead innovation, scale their operations, and offer a differentiated and sustainable user experience.

At Inetum, we continuously monitor the evolution of TM Forum standards (DT4DI, AN) to evaluate their potential return on investment from a business perspective. We have extensive experience in digital transformation projects for Telco operators, which allows us to understand the needs and challenges of the sector. We are currently working on defining conceptual models that will allow us to explore these approaches in future test labs or PoCs. Our goal is to lay the necessary foundations to demonstrate, in a controlled environment, how digital twins and AI can generate real value—both in operational efficiency and in new revenue opportunities—for our clients in the Telco sector.

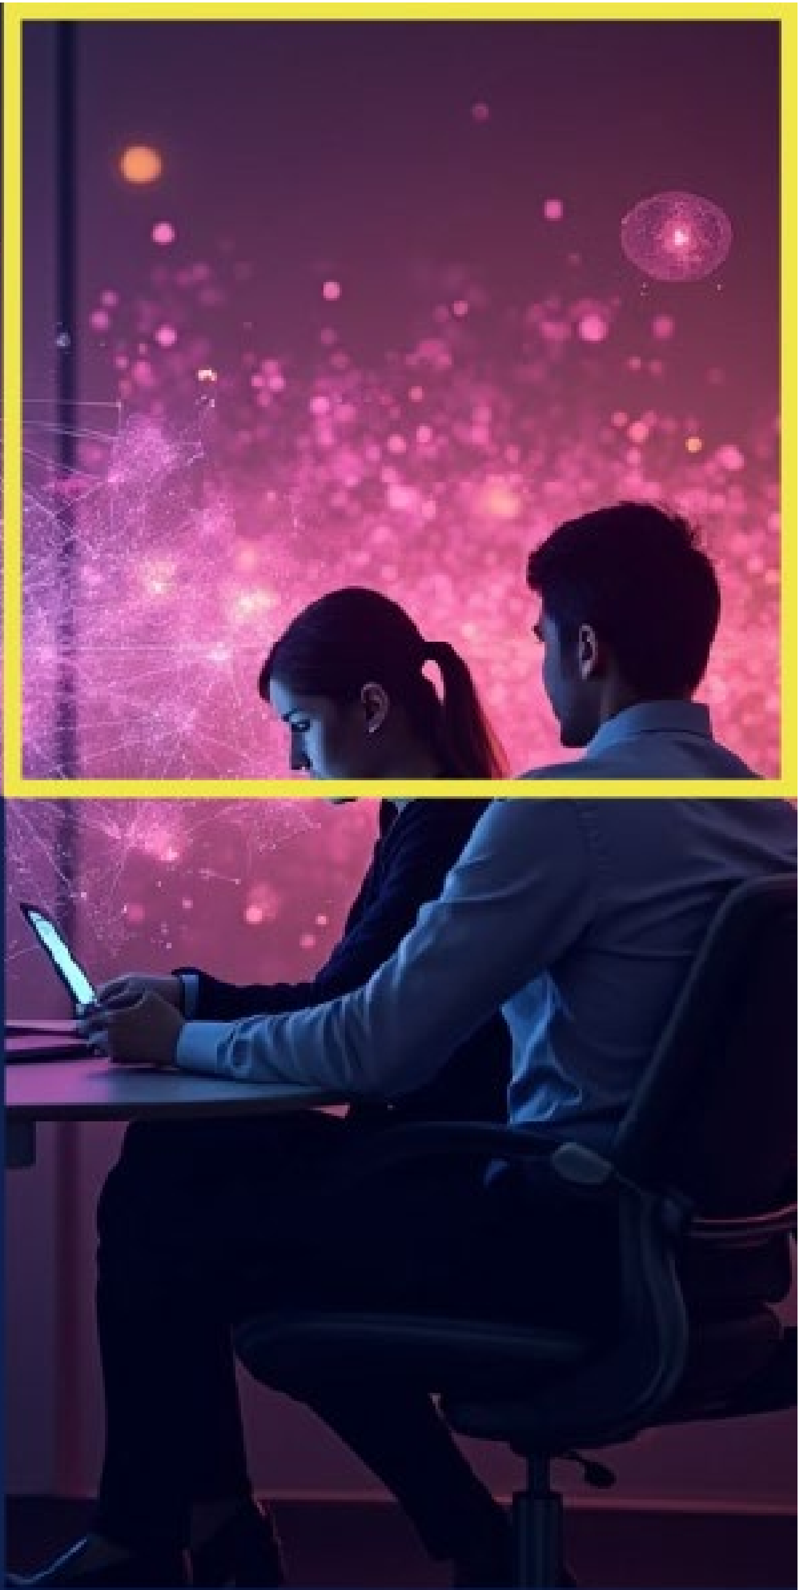

09

Contribution and Contact Team

9 Contribution and Contact Team

This White Paper has been developed by the **Inetum Consulting** team, with the collaboration of experts in **digital transformation, automation, and data analytics for the Telco sector**. For more information or inquiries, you can contact the following specialists:

Victoria Pérez Carrasco
Director of Telecommunications and Energy
victoria.carrasco@inetum.com

José Pardo Vázquez
Senior Manager in Telco-Media
jose.pardo@inetum.com

Osmar Enrique Roman Hidalgo
Senior Consultant in Telco-Media
osmar.roman @inetum.com

Dalia Vásquez Leyva
Senior Consultant in Telco-Media
dalia.vasquez@inetum.com

Santiago Fernández Tinoco
Trainee in Telecommunications Sector
santiago.fernandez.ext@inetum.com

We thank the entire **Inetum Consulting** team for their contribution to the development of this document and for their commitment to the evolution of the telecommunications sector.

Inetum Consulting

We transform the telecommunications industry through the combination of agile processes, automation, advanced analytics, and a deep understanding of the technologies that enable the future of connectivity.

inetum.

inetum.com