Data-Driven Redesign of a Logistics Process in a Telecommunications Company: From Diagnosis to Transformation

Inês Afonso Furtado

Master's Dissertation

Supervisor at FEUP: Prof. Maria João Pires

U.PORTO
FEUP FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO

Master in Industrial Engineering and Management

2025-06-27

Abstract

In a sector such as telecommunications, characterized by high operational complexity and increasing pressure to ensure agility, service quality, and customer satisfaction, effective inventory management plays a key role. Equipment traceability, data integrity, and the ability to quickly detect and resolve stock inconsistencies — understood as discrepancies between physical stock and the quantities recorded in the system — are critical factors for the proper functioning of the supply chain and, consequently, for organizational success.

In this context, the present dissertation focuses on improving a critical stock reconciliation process within a Portuguese telecommunications company.

The process under analysis exhibited several weaknesses: lack of performance indicators and monitoring mechanisms, the need for coordination across multiple systems and interfaces, heavy dependence on a single operator, poorly structured communication among stakeholders, and activities of questionable added value — all of which contributed to a process that was opaque, slow, and difficult to manage.

To address these limitations, the project pursued two main objectives: to increase the visibility of the reconciliation process at the management level and to enhance its operational efficiency, making it more streamlined and effective. To this end, the work was carried out across several complementary stages.

The project began with the framing of the problem and the alignment of objectives among stakeholders, ensuring a shared understanding of the project's priorities and goals. This was followed by a detailed analysis of the current process (AS-IS) and the consolidation of relevant data, involving the structured treatment and integration of information collected through extraction, transformation, and loading (ETL) processes. The next stage involved the development of a dashboard in Power BI and the redesign of the process, which included the elimination of unnecessary tasks, the redistribution of responsibilities, and the introduction of automation, with the aim of reducing manual intervention and accelerating the resolution of inconsistencies.

With the new solution defined, a pilot test was carried out, during which the redesigned process was implemented and monitored in a real environment, followed by an evaluation of its performance.

Among the main achievements, the visibility provided by the dashboard stands out, allowing for the first time the systematic monitoring of indicators, historical performance analysis, and the creation of a structured foundation for future analyses, thereby strengthening decision-making. Regarding the redesigned process, it is estimated that the weekly time savings amount to 6 hours and 10 minutes, equivalent to an annual gain of approximately 320 hours and 38 minutes, representing a cost saving of around 3 206.40 €. These results highlight the project's contribution to improving operational efficiency and reinforce the importance of adopting data-driven management practices.

Resumo

Num setor como o das telecomunicações, caracterizado por elevada complexidade operacional e crescente pressão para garantir agilidade, qualidade de serviço e satisfação do cliente, a gestão eficaz de inventário assume papel fundamental. A rastreabilidade de equipamentos, a integridade dos dados e a capacidade de detetar e resolver rapidamente inconsistências de stock — entendidas como discrepâncias entre o stock físico e o registado no sistema — são fatores críticos para o bom funcionamento da cadeia de abastecimento e, consequentemente, para o sucesso organizacional.

Neste contexto, a presente dissertação incide sobre a melhoria de um processo crítico de regularização de inconsistências de stock numa empresa portuguesa do setor das telecomunicações.

O processo em análise apresentava diversas fragilidades: ausência de indicadores de desempenho e de mecanismos de monitorização, necessidade de coordenação entre múltiplas interfaces e sistemas, forte dependência de um único operador, comunicação pouco estruturada entre os intervenientes e a existência de atividades cuja utilidade efetiva era questionável — fatores que resultavam num processo opaco, moroso e difícil de gerir.

Com o objetivo de ultrapassar estas limitações, o projeto desenvolvido teve duas finalidades principais: aumentar a visibilidade do processo de regularização junto das camadas de gestão e promover a sua melhoria operacional, tornando-o mais eficiente. Para isso, o trabalho foi desenvolvido em várias etapas complementares.

Inicialmente, procedeu-se ao enquadramento do problema e ao alinhamento de objectivos entre as partes interessadas, garantindo um entendimento comum das prioridades e metas do projeto. Seguiu-se uma análise detalhada do processo em vigor (*AS-IS*) e a consolidação dos dados relevantes, com tratamento e integração estruturada da informação recolhida através de processos de extração, transformação e carregamento (*ETL*). A etapa seguinte consistiu no desenvolvimento de uma *dashboard* em *Power BI* e no redesenho do processo, o qual contemplou a eliminação de tarefas desnecessárias, a redistribuição de responsabilidades e a introdução de automatismos, com o objectivo de reduzir a intervenção manual e acelerar a resolução de inconsistências. Com a nova solução delineada, avançou-se para a condução de um teste-piloto, em que o processo reformulado foi monitorizado em ambiente real, tendo-se procedido posteriormente à sua avaliação.

Entre os principais ganhos, destaca-se a visibilidade proporcionada pela *dashboard*, que permitiu, pela primeira vez, a monitorização sistemática de indicadores, análise histórica de desempenho e uma base estruturada para análises futuras, fortalecendo a tomada de decisão. Em relação ao processo redesenhado, estima-se uma poupança semanal de 6 horas e 10 minutos, o que equivale a uma economia anual de aproximadamente 320 horas e 38 minutos, correspondendo a cerca de 3 206.40 €. Estes resultados evidenciam a contribuição do projeto para a melhoria da eficiência operacional e reforçam a importância da adoção de práticas de gestão *data-driven*.

Acknowledgments

Life can be seen as a collection of chapters, each defined by the people we meet, the lessons we learn, and the moments that shape us. As I prepare to close another one, I want to take this space to thank everyone who, in their own way, has contributed to this five-year journey and made it truly meaningful.

First, to NOS Comunicações S.A., thank you for the opportunity to carry out this work. In particular, to Sérgio Santos and Ana Rocha, my company supervisors, thank you for your support, availability, and guidance throughout the project.

To my academic advisor, Professor Maria João Pires, my heartfelt thanks for your attentiveness, encouragement, and understanding through all the ups and downs.

To my parents, for the values you instilled in me and your unwavering belief in my path, especially on the days when doubt spoke louder. To my mother in particular, for being my greatest support.

To Bernardo, for his companionship, encouragement, and above all, his heart.

To the city of Guimarães, for being such an unforgettable chapter in my life. To the friends I made there and the irreplaceable experiences I will never forget.

To all my friends, both those who have been with me since day one and those who joined along the way, thank you for your loyalty, friendship, and the fresh breeze you bring into my life.

To my parents,

Contents

Acronyms and Symbols

BI	Business Intelligence
BPM	Business Process Management
BPR	Business Process Reengineering
CSF	Critical Success Factor
DAX	Data Analysis Expressions
ERP	Enterprise Resource Planning
ETL	Extract, Transform, Load
IRI	Inventory Record Inaccuracy
IT	Information Technology
ITSM	Information Technology Service Management
KPI	Key Performance Indicator
OLAP	Online Analytical Processing
PM	Performance Measurement
SCM	Supply Chain Management
SKU	Stock Keeping Unit
SLA	Service Level Agreement
SP	Service Provider
SSBI	Self-Service Business Intelligence
UML	Unified Modeling Language
VBA	Visual Basic for Applications
WO	Work Order

List of Figures

List of Tables

Chapter 1

Introduction

This dissertation was developed within the Master's Degree in Industrial Engineering and Management at the Faculty of Engineering of the University of Porto and investigates the redesign of a key process for resolving inventory inconsistencies at a Portuguese telecommunications company. Section 1.1 introduces the company and the supply chain team involved in the initiative. Section 1.2 outlines the motivations behind the project, emphasising its relevance both at a broader organisational level and within the company's internal context. Sections 1.3 and 1.4 define the main goals and describe the methodological approach adopted, respectively. Finally, Section 1.5 concludes the chapter with an overview of the structure of the dissertation.

1.1 Supply Chain at Nos Comunicações S.A.

NOS was founded in 2014 following the merger between ZON Multimédia and OPTIMUS Telecomunicações. ZON Multimédia was renowned for being a market leader in cable television, broadband internet, and cinema operations, while Optimus Telecomunicações excelled as a mobile telecommunications operator. This merger has made it possible to obtain financial synergies, improve the quality of services and extend the range of offers at more competitive prices.

The company currently operates through several subsidiaries, with NOS Comunicações, S.A. being the most prominent, reinforcing the brand's name and providing internet, voice, and data solutions for both residential and business customers. Beyond telecommunications, NOS also holds a leading position in the entertainment sector through NOS Lusomundo Cinemas, S.A., which operates Portugal's largest cinema chain, and NOS Audiovisuais, S.A., which specializes in film distribution.

NOS has also earned widespread recognition for its excellence in service. It was awarded the distinction of operating Europe's fastest 5G network in 2023 and Portugal's fastest in 2024. That same year, it was named "Recommended Brand 2024" after achieving top customer satisfaction scores in television, internet, and fixed-line services. The organisation also leads in innovation, registering the highest number of patent submissions in Europe in 2024, and actively supporting the national innovation ecosystem through its 5G Hub program, which backs startups and university-driven technological initiatives.

This project was developed at NOS Comunicações, S.A., hereinafter referred to simply as NOS, in close collaboration with the Fixed Equipment Planning team, which is part of the Procurement and Supply Chain department. This team manages the entire logistics lifecycle of fixed devices—such as set-top boxes, routers, signal extenders, and landline phones—from procurement and distribution to returns and refurbishment. The team is composed of five members, represented by the letters A, B, C, D, and E, each with specific responsibilities. The operational workflow is depicted in Figure 1.1.

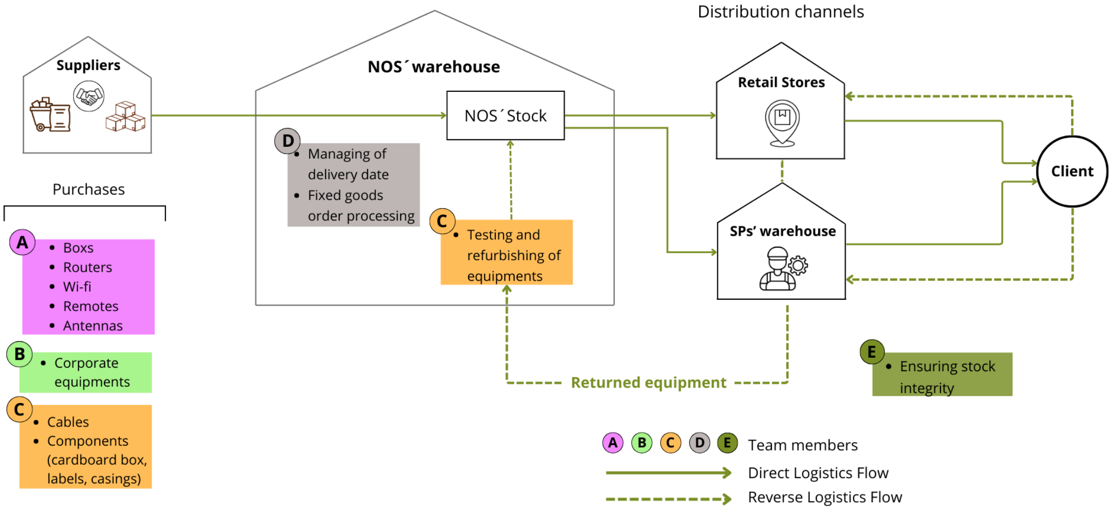

Figure 1.1: Fixed Equipment Planning team overview

Procurement responsibilities are divided among three key roles: A acquires equipment for residential customers, B manages procurement for business clients, and C oversees the purchase of components for refurbishment. Once secured, D coordinates deliveries with suppliers, ensuring proper reception and quality verification at NOS warehouse.

Stock replenishment, also managed by B, ensures adequate inventory across distribution channels. The main channels include NOS stores, which have become strategic points enabling a "pick and go" approach when installing and replacing equipment can be easily handled by the customer; and third-party companies called Service Providers (SPs), responsible for installation and maintenance at customer residences. Supply to these channels follows the stock management system, *Modelo Integrado de Planeamento*, developed by LTPlabs, which forecasts stock needs based on historical consumption and demand predictions. Faulty or replaced equipment must be returned by clients through stores or SPs, initiating the reverse logistics process, coordinated by C. After assessment, devices are either repaired and refurbished or discarded when necessary. Refurbished equipment are repackaged and reintegrated into stock, reducing costs and optimizing resources.

Operator E, although not directly involved in the execution of the supply chain, is responsible for ensuring stock integrity, enabling effective tracking of all equipment and encouraging the resolution of potential stock inconsistencies associated with both direct and reverse logistics flows. This dissertation examines one of its core processes with the aim of assessing its performance and identifying opportunities for improvement.

1.2 Motivation for the Project

Ensuring stock integrity remains one of the most significant challenges in today's supply chains, particularly in sectors such as telecommunications, where accurate records of physical assets are essential for maintaining operational efficiency, service continuity and customer satisfaction (Doe, 2022).

The challenge becomes even more complex when multiple stakeholders are involved in the supply chain. The use of different systems, independent teams, and varying procedures increases the likelihood of mismatches between physical stock and system records, making inventory control more difficult and less reliable. This type of issue is recognised in the literature as Inventory Record Inaccuracy (IRI). Some definitions, such as the one proposed by Sheppard and Brown (1993), consider a record inaccurate when the discrepancy between the recorded and actual inventory quantity exceeds 1% of the recorded balance.

Despite ongoing improvements in logistics management and information systems, IRI remains a widespread and persistent challenge. For instance, DeHoratius and Raman (2008) found that IRI can affect over 65% of inventory records in retail and warehousing settings. The consequences of this problem are well documented and have a significant impact on several aspects of organisational performance.

- **Operationally**, these inaccuracies can lead to stockouts that go unnoticed; delays in restocking or processing returns and frequent manual reconciliation tasks that consume valuable time and resources. Moreover, and due to low confidence in system data, companies often rely on internal audits, gap scans, and additional physical counts. These corrective efforts can account for up to 10% of total labour costs in logistics operations (Rekik et al., 2019).

- **Financially**, even small errors can have significant consequences. In one case study conducted in Indonesia, an average inventory error of just 0.53% led to estimated losses of IDR 154 million, showing how minor inaccuracies can distort stock valuation and compromise the reliability of financial reports (Linuwih and Handayati, 2025).

- **Strategically**, IRI reduces trust in the data used for planning, forecasting, and decision-making. It can hinder internal coordination, pose challenges for compliance and auditing, and negatively affect the organisation's credibility with customers, partners, and regulators.

NOS, aware of these implications, has established processes and dedicated teams for stock regularisation, with particular focus on two key points in the supply chain: retail stores and SPs. Although these entities operate with some autonomy, ensuring the availability of the equipment to the customer and its subsequent collection at the end of its useful life, all stock remains the property of NOS and is tracked through its Enterprise Resource Planning system (ERP) — SAP ERP — which provides real-time visibility into the location and financial status of each item. For this visibility to be meaningful, system data must accurately reflect physical reality.

Figure 1.2 presents a comparison of the main types of stock inconsistencies that fall within the scope of the Fixed Equipment Planning team. It outlines the logistics flow in which these discrepancies occur, describes the nature of the errors, and identifies the entities responsible for their detection and reporting. Any other types of stock inconsistencies that may arise within the direct or reverse logistics flows fall outside the scope of this team and are assigned to another team within NOS.

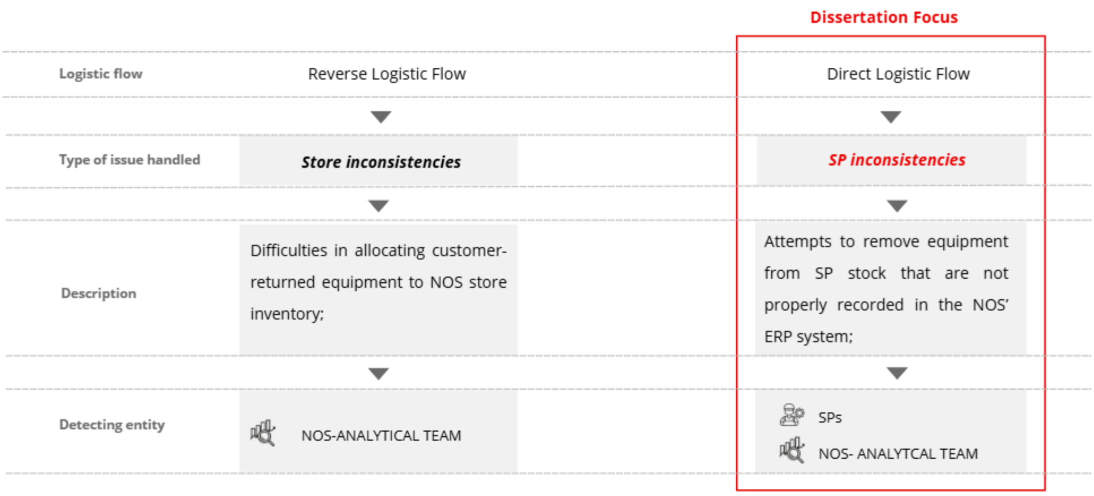

Figure 1.2: Types of stock inconsistencies handled by the Fixed Equipment Planning team

Currently, the handling of inconsistencies that arise in retail stores, *Stores Inconsistencies*, does not represent a critical issue, as the process is fully automated. It is supported by SAP queries integrated with VBA (Visual Basics for Applications) scripts, which enable fast, standardized, and reliable resolution.

In contrast, *SP Inconsistencies* are still handled manually, with tasks fragmented across several platforms, a factor that complicates automation. A further challenge is the lack of managerial visibility into this activity, including how it operates, the effort it entails, and the constraints it imposes, all of which raise concerns about its effectiveness and long-term sustainability. Clarifying and optimising this operational flow aligns closely with the mission of the Fixed Equipment Planning team, which has been actively investing in the continuous improvement of internal processes.

1.3 Project Goals

The main objective of this dissertation is to analyse and optimise the *SP Inconsistencies* regularisation process, strengthening the decision-making capacity of management teams and, consequently, contributing to greater operational efficiency. To this end, the project is structured around the following specific objectives:

- **Diagnose** the current process by identifying structural limitations and operational inefficiencies that hinder performance;
- **Define and implement** a consistent and scalable data infrastructure to enable reliable monitoring and analysis;
- **Develop** a set of Key Performance Indicators (KPIs) aligned with operational priorities and strategic goals;
- **Design and build** a dashboard to ensure real-time visibility, enhance transparency, and support data-driven decision-making;
- **Redesign** the existing stock regularisation workflow to eliminate non-value-adding activities, automate routine tasks, and decentralise responsibilities;
- **Pilot and evaluate** the proposed improvements through a controlled implementation, using objective performance metrics to compare pre- and post-intervention results, including operational gains and estimated cost savings.

1.4 Methodology

To conduct the project, the Action Research methodology was adopted, a practice-oriented approach particularly suited to applied research in organisational contexts (Melin and Axelsson, 2016). This methodology involves the researcher directly in the problem-solving process and follows the principle of *learn by doing*, combining action and reflection in iterative cycles of improvement (O'Brien, 1998).

Over a five-month period, the project was developed through five iterative stages: (1) planning, (2) data collection, (3) analysis and action planning, (4) implementation, and (5) evaluation. Figure 1.3 provides a visual summary of these stages, adapted to the specific context and workflow of the project.

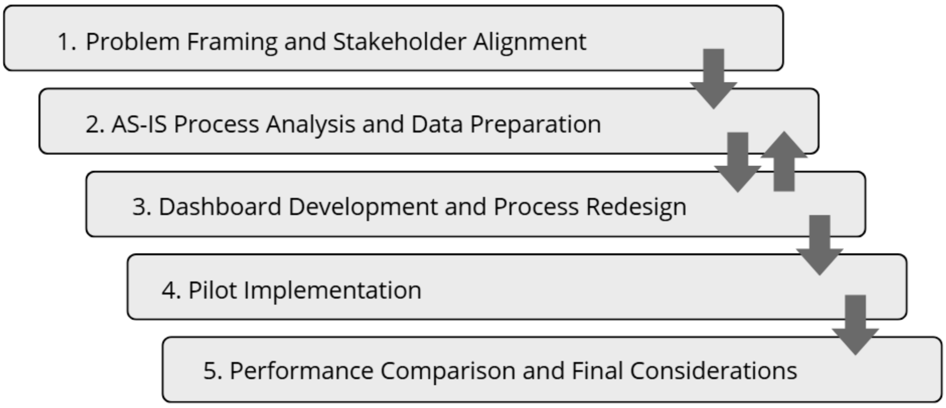

Figure 1.3: Project development stages based on the Action Research methodology

The initial phase focused on understanding the problem within its organisational context. This included defining the scope of the intervention, identifying relevant stakeholders, and clarifying the project's objectives. A preliminary literature review, supported by exploratory meetings, helped identify the main challenges and provided a clearer understanding of the problem's relevance and the need for intervention.

The second phase focused on gathering information. The current process (AS-IS) was initially modeled through direct observation and the help of informal interviews. Although some inefficiencies were identified at this stage, a robust, data-driven foundation was required to validate and quantify performance issues. To achieve this, quantitative data from multiple internal sources were consolidated using an Extract, Transform and Load (ETL) pipeline, a procedure that extracts data from various sources, converts it into a suitable format, and loads it into a target database or data warehouse. The pipeline was developed in Python and iteratively improved as new requirements emerged.

With reliable data in place, the third stage involved designing a Power BI dashboard to analyse process performance and support future monitoring. This tool was refined based on stakeholders feedback to ensure that it accurately reflected the process reality and supported decision-making. Insights obtained from the dashboard, together with the previously mapped workflow, guided the redesign of the regularisation process (TO-BE), with a focus on eliminating inefficiencies and reinforcing control mechanisms.

The fourth stage consisted of piloting the redesigned process to assess its behaviour under real operational conditions.

In the final stage, the pilot results were evaluated and compared with the previous process. Qualitative feedback from stakeholders was also gathered. The insights gained were documented and used to inform a set of final recommendations.

1.5 Dissertation Structure

The dissertation is organized into six chapters. The first chapter frames the project by introducing the company and the team where it was developed, and highlighting the importance of resolving stock inconsistencies. It also defines the project's goals and outlines the adopted methodology.

The second chapter explores the key theoretical concepts that support the project. It covers topics such as process improvement and modelling, the role of data in decision-making, the various stages of a Business Intelligence (BI) architecture, and the importance of performance measurement through the definition of KPIs.

The third chapter sets the context of the problem, beginning with a brief overview of the situation under analysis. It identifies the main stakeholders and systems involved, highlights the sources that generate stock inconsistencies, and provides a detailed analysis of the current resolution process. The chapter concludes by identifying the main improvement opportunities, grouping them into two thematic areas: Information Systems and Process and Operations.

The fourth chapter presents the structured approach adopted to address the previously identified improvement opportunities. It is organised into three main pillars: first, the consolidation and organisation of data sources to support the development of an effective monitoring tool; second, the definition of the most relevant KPIs and their translation into a visual dashboard; and finally, the redesign of the existing process for resolving *SP inconsistencies*.

The fifth chapter discusses the main outcomes and findings from the project's implementation. It is divided into two sections: the first highlights insights obtained from the dashboard analysis; the second evaluates the effectiveness and estimated gains of the redesigned workflow.

The final chapter summarises the project's key findings and contributions, highlighting the improvements achieved, the challenges encountered during execution, and the proposed directions for future development.

Chapter 2

Literature Review

This chapter presents a comprehensive review of the theoretical and conceptual foundations relevant to the development of this project. The selected topics reflect the multidimensional nature of the problem addressed, encompassing process improvement, data reliability, and performance measurement.

Section 2.1 examines the main existing approaches to business process improvement. Section 2.2 underscores the importance of process mapping for understanding workflows, identifying inefficiencies, and guiding improvement initiatives. Section 2.3 introduces the concept of a data-driven culture, emphasising its relevance in supporting informed, evidence-based decision-making. Section 2.4 explores the BI architecture, with particular focus on key components such as the ETL process and the use of dashboards for effective data visualisation. Finally, Section 2.5 highlights the importance of performance measurement, with particular attention to the role of KPIs, expanding the theoretical framework with concrete examples directly related to the problem under analysis.

2.1 Managing and Redesigning Processes

An initial assessment of a business process is crucial to determine whether it should be preserved with minor improvements, fundamentally reengineered, or completely discontinued. Depending on the degree of change required, distinct structured methodologies can be applied. In this context, Business Process Management (BPM) and Business Process Reengineering (BPR) emerge as two fundamental methodologies for improving organisational processes. BPM is defined as "a collection of technologies capable of translating process models into computer-supported activities, relieving organisational agents of routine tasks" (Antunes and Mourão, 2011). In other words, BPM enables organisations to model, monitor, and continuously improve operational processes through the integration of people, applications, documents, and digital workflow. The primary objective is to analyze the current process and propose future improvements. In contrast, BRP involves a more disruptive change. Hammer and Champy (1993) describe it as the "fundamental rethinking and radical redesign of business processes to achieve dramatic improvements in critical performance measures such as cost, quality, service, and speed”. BPR gained popularity in the 1990s following the publication of "Reengineering Work: Don’t Automate, obliterate”, with impactful results: Ford reduced its accounts payable workforce by 75%, Mutual Benefit Life improved insurance underwriting efficiency by 40%, and Xerox reduced inventory cycle times by 70% (Anand et al., 2013).

The choice between BPM and BPR should be guided by the degree of transformation needed. If the goal is to enhance existing processes through incremental improvements, BPM is the appropriate path. Conversely, if processes are ineffective, outdated, or misaligned with organisational strategy, and radical change is required to achieve significant performance gains, then BPR is the suitable approach.

2.2 Process Mapping

Process mapping is a key technique in organisational process management, designed to visually represent the steps, roles, and flows involved in a given workflow. By making process structures explicit, it enables the identification of inefficiencies, bottlenecks, redundancies, and other performance gaps, thereby supporting informed decision-making and operational improvement (Eyeregba et al., 2024). Mapping the current state of a process — the AS-IS — fosters transparency and shared understanding across stakeholders. This shared view forms the basis for identifying improvement opportunities and for designing a more efficient and value-driven future state — the TO-BE.

In practice, process mapping efforts tend to concentrate on core operational processes — those that contribute directly to value creation and customer delivery — while support and management processes are typically represented as enablers of these value streams (Harmon, 2010; Dijkman et al., 2011). Such prioritization ensures that improvement initiatives are focused on the areas with the greatest potential impact on organisational performance.

Among the various tools used for this purpose, the Swimlane Diagram is particularly effective. It structures the process into horizontal or vertical lanes, each representing a department, function, or role, and clarifies the division of responsibilities and interdependencies. This visualisation not only aids in process analysis but also serves as a practical reference for defining work procedures and even onboarding new staff.

To remain relevant, process mapping should be integrated into a continuous improvement cycle. Regular updates are therefore essential to reflect changes in technology, operations, or strategic direction, ensuring that processes evolve in alignment with organisational goals.

2.3 Data-driven Culture

The business environment is constantly transforming, becoming increasingly dynamic and complex. This evolution, driven by globalization, market instability, and technological advances, has led to exponential growth in data volume, velocity, and variety. In this scenario, effective information management becomes essential to maintain long-term competitiveness and sustainability of companies (Fernandes et al., 2023; Rausch et al., 2013).

It can therefore be asserted that data assumes a central role in society, being considered a strategic asset and one of the main drivers of organisational leverage, promoting more informed and decision-making (Fernandes et al., 2023). In fact, organisations that embrace data-driven approaches tend to outperform their industry peers, achieving around 5% higher productivity and 6% greater profitability (Brynjolfsson et al., 2011).

This strategic role of data becomes particularly evident in the diagnosis and optimization of organisational processes. While process modeling offers valuable structural insights, highlighting critical points, an effective diagnosis requires data that reflects the actual behavior of operations. Without the systematic collection and analysis of operational data, any attempt to evaluate or improve a process remains partial and susceptible to bias.

Despite growing recognition of its importance, the concept of a data-driven culture still lacks a uniform definition in both theory and practice, resulting in varied interpretations. Anton et al. (2023) offer a comprehensive perspective, arguing that technology adoption and information access alone are insufficient. To cultivate a truly data-driven culture, organisations must commit strategically and culturally to integrating data into decision-making processes. These authors advocate for investment not only in tangible elements—such as technological infrastructure and structured processes—but also in intangible aspects like data literacy and organisational mindset. In this context, Business Intelligence systems have become key enablers of data-driven decision-making. These tools transform raw data into structured and actionable knowledge, allowing organisations to discover new business opportunities, improve operational efficiency, reduce costs, and ultimately make more informed strategic decisions (Costa, 2024).

2.4 BI Concept and Architecture

Business Intelligence is a dynamically evolving construct that lacks a universally standardized definition within academic literature. Its origins can be traced back to 1958, when IBM researcher Hans Peter Luhn introduced the concept in his seminal article "A Business Intelligence System", describing an automated framework to disseminate relevant information efficiently across organisations (Luhn, 1958). In its contemporary understanding, BI is commonly defined as "an integrated, company-specific approach, based on Information Technology (IT) and aimed at supporting managerial decision-making" (Rausch et al., 2013). Therefore, BI today represents a structured, technology-driven process designed to collect, analyze, monitor, and forecast business scenarios, continuously evolving to meet new demands (Fernandes et al., 2023; Zhao, 2021).

Rather than being conceptualized as a standalone tool or technology, BI should be understood as an integrated, structured process composed of six fundamental phases. The first phase consists of defining a BI strategy aligned with the organisation's objectives. This step involves identifying business goals and defining key indicators essential for monitoring the performance of adopted initiatives.

Next occurs the ETL process which plays a central role in consolidating data from multiple sources, applying business rules for transformation, and loading the results into destination repositories. This is followed by the data analysis stage, in which advanced techniques—such as data mining and OLAP (Online Analytical Processing)—allow for the generation of relevant information for decision-making. Finally, the results obtained are presented to stakeholders through structured reports, interactive dashboards, scorecards, or other visualisation tools, ensuring clear and accessible presentation of information. (Costa and Santos, 2012; Foley, 2010).

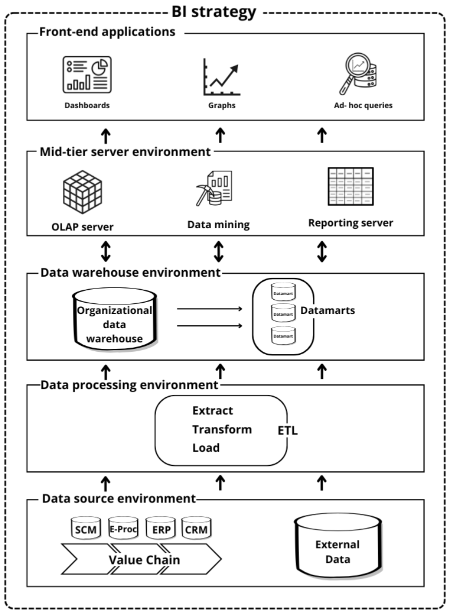

Figure 2.1: BI Architecture. Adapted from: (Costa and Santos, 2012; Rausch et al., 2013)

2.4.1 ETL Process

The ETL process is one of the most important components of modern BI architecture. Data quality depends heavily on the performance of the ETL process, which-if poorly designed-carries the risk of reproducing the garbage received from various data sources, compromising future analysis (Garbage In, Garbage Out - GI-GO) (Hamed and Ghozzi, 2015)

At its core, the ETL process follows a three-stage sequence:

- **Extract:** Data is gathered from a variety of sources with varying levels of complexity depending on the data source. This can encompass both internal and external sources. Among internal sources, enterprise systems such as Enterprise Resource Planning (ERP) and Supply Chain Management (SCM) stand out, while examples of external sources include social media and market data. The collected data can present different levels of structuring: structured, when organized in tabular format; semi-structured, such as JSON and XML files; and unstructured, such as images, videos, and free text. There are two approaches to data extraction: static ingestion, which involves capturing all available records from a source and passing them on to the next steps of the ETL process; and incremental ingestion, in which only new, updated, or recently deleted records are retrieved during each iteration.

- **Transform:** This is the core and most complex phase, where data is transformed. It includes tasks such as filtering, standardization, null handling, and semantic mapping, often automated through system processes. Business rules are applied during this phase to ensure consistency, and the approach always varies according to the specific needs of the data.

- **Load:** Transformed data is then inserted into the target repository-Data warehouses. Alternatively, they can be organized in Data Marts, subsets of the Data Warehouse segmented by business unit or specific area, allowing for more agile and targeted queries (Foley, 2010). There are two main approaches to loading data into a Data Warehouse. The refreshing method involves replacing all existing data with a completely new version and is typically used during the initial setup of the Data Warehouse. The updating method, on the other hand, adds or updates only the data that has changed, making it suitable for ongoing maintenance (Hamed and Ghozzi, 2015).

The ETL process is, therefore, fundamental for ensuring data consistency, making it suitable for analysis, and addressing the initial heterogeneity and disorder present in raw data sources. However, building and maintaining an efficient ETL process is widely recognized in the literature as a complex, time-consuming, and error-prone endeavor. It requires interdisciplinary collaboration among business analysts, database architects, and developers (Hamed and Ghozzi, 2015). One of the main key challenges is to reconcile evolving business requirements with the structural and semantic diversity of multiple data sources. Moreover, and in order to remain resilient and adaptable, ETL processes must support what-if analysis and allow for continuous refinement.

2.4.1.1 Data and ETL Pipeline

Traditionally, ETL processes were executed manually and in a segmented fashion, requiring human intervention at various stages. This approach presented clear limitations in terms of scalability, reliability, and operational efficiency. With the evolution of data technologies and the growing demand for more robust and automated solutions, the concept of the ETL pipeline emerged, representing the full automation of the ETL process. An ETL pipeline is part of a broader concept known as a data pipeline. According to the literature, a data pipeline is defined as an automated sequence of operations that manages the flow of data from its source to a final destination—typically a Data Warehouse. Each stage consumes the output of the previous one and processes it according to predefined transformation rules, ensuring consistent, efficient, and reliable data flow throughout the pipeline (Alstyne et al., 2016). The automation of the ETL process significantly reduces the risk of human error, improves data quality, and enhances overall data management effectiveness (Ogunsola et al., 2022). In addition, it allows for more agile responses to demands for data updates, integration, and continuous monitoring.

2.4.2 Data Visualisation

Given the way the human brain processes information, identifying trends and patterns in texts or numerical data can be a complex and time-consuming task. Naturally, the brain is more receptive to images and visual representations (Lousa et al., 2019).

According to Zheng (2018) , "data visualisation is the graphical or visual method of presenting data." This presentation can take various forms, such as text, tables, or graphs. The selection of the most suitable format depends on several factors, including the nature of the data, the target audience, and the specific objectives of the analysis.

Regardless of the chosen format, effective presentations must provide visual cues that direct user attention to discrepancies, patterns, and outliers, enabling an intuitive interpretation of the information. In this way, even decision makers with limited data literacy can use their cognitive and visual skills to identify critical points requiring more in-depth analysis, reducing reliance on speculation and subjective interpretations (Zheng, 2018; Solanki et al., 2024).

To ensure that data is accessible and comprehensible to a broad audience, it is essential to structure it in a coherent and logical narrative. As Herschel and Clements (2017) emphasize, what seems self-evident to an analyst may not necessarily be clear to their audience. This underscores the need for a well-designed visual structure that not only represents data but also enhances its interpretation within the specific context of the analysis. The integration of data visualisation with narrative techniques emerges as a powerful method for improving comprehension and reducing ambiguity. Herschel and Clements (2017) further argue that “stories are essential, as they not only inform the audience but also inspire action."

The growing significance of visual information representation has positioned Self-Service Business Intelligence (SSBI) as a critical approach for decentralizing data analysis. This methodology empowers non-technical users to independently access, interpret, and explore data without relying on IT teams for support. By democratizing access to information, SSBI facilitates analyses that address both strategic and operational concerns, expediting decision-making and enabling business units to independently extract value from data (Lousa et al., 2019).

Among the various SSBI tools available, Tableau and Power BI stand out as two of the most widely adopted options. Tableau is particularly well-suited for scenarios involving complex and high-volume datasets. However, its high licensing costs may pose a limitation for some organisations (Lousa et al., 2019). In contrast, Power BI is often the preferred choice for organisations operating with tighter budgets. Recognized as a leader in data visualisation by the *Gartner Magic Quadrant for Analytics and BI Platforms 2024* (Microsoft, 2024), Power BI offers a user-friendly interface, seamless integration with multiple data sources, and robust functionalities, including Power Query, Power Pivot, Power View, and Power Map (Lyon, 2019).

Furthermore, in the context of ETL, Power BI simplifies data transformation and preparation, eliminating the need for intricate formulas. Its data modeling capabilities are further enhanced by the inclusion of DAX (Data Analysis Expressions), enabling efficient and advanced calculations (Santo Andrade, 2020).

Ultimately, the selection of the most appropriate tool for an organisation depends on factors such as budget constraints, the skill levels of professionals, and the specific goals to be achieved.

2.4.2.1 Dashboards

There are several definitions that describe the fundamental characteristics of a dashboard. According to Few (2006), it is "a visual display of the most important information needed to achieve one or more objectives; consolidated and arranged on a single screen so the information can be monitored at a glance." Kitchin et al. (2015) adds that a dashboard "should act as a translator and not simply as a mirror, defining the forms and parameters by which data are communicated, and therefore what the user can see and interact with."

To ensure that each user accesses the data relevant to them, while avoiding information overload, dashboards can be divided and classified into three categories: strategic, tactical, and operational. Strategic dashboards are intended for executives and allow them to monitor the implementation of strategic goals, organisational performance, and support long-term planning. Tactical dashboards, on the other hand, are aimed at middle managers and contribute to controlling departmental processes or specific projects, facilitating decision-making. Finally, operational dashboards provide real-time information to frontline staff, enabling the monitoring of daily activities and the quick identification of problems (Rahman et al., 2017).

Regardless of the type of dashboard being developed, there are cross-cutting aspects that must always be considered in advance: a dashboard should not overwhelm users or present an excessive volume of data; the selection of KPIs to be analyzed must be rigorous; the dashboard should be aligned with existing workflows, naturally integrating into the user's routines and processes; it should support scenario analysis and exploration, incorporating functionalities such as *drill-down* and *drill-through*, without, however, creating excessive dependency on these tools; its design should allow all key information to be displayed on a single screen, eliminating the need to scroll; and the choice and combination of colors should be appropriate, enabling effective data encoding and the assignment of meaning. In general, a dashboard should be based both on functional characteristics – which define what it is capable of doing – and on visual characteristics – which determine how the information is presented (Bach et al., 2022; Miguéis, 2024; Rahman et al., 2017).

These principles have been successfully applied in real-world settings. A notable example is provided by NetApp. To address inefficiencies caused by fragmented data access, NetApp implemented a centralized BI environment structured around role-based dashboards — strategic, tactical, and operational — aligned with different user needs. A core design principle was to ensure that any critical information could be accessed within three clicks. As a result, the company accelerated decision-making, improved visibility across departments, and significantly reduced manual data consolidation efforts (Eckerson, 2011).

2.5 Performance Measurement and KPIs

Performance measurement (PM) is a systematic evaluative process used by organisations to monitor and assess their outputs, outcomes, and impacts. It functions as a vital management tool, enabling stakeholders to make evidence-based decisions and implement timely corrective actions to enhance program effectiveness (Nkwake, 2023). Of the various available approaches, KPIs remain one of the most widely adopted methods for evaluating performance.

By definition, KPIs are measurable values that indicate how effectively an organisation is achieving its key objectives (Kantor et al., 2019). However, it is important to note that not all performance metrics qualify as KPIs. KPIs represent a strategically selected subset of indicators that provide real added value by guiding decision-making processes—rather than merely monitoring routine activities. As noted by Franceschini et al. (2007) , the most critical step in designing a PM system is not to identify as many indicators as possible, but to focus on those that truly reflect the value-creating areas of the organisation.

In the literature, KPIs are commonly classified into three main categories (Badawy et al., 2016). *Lagging indicators* assess past performance, allowing organisations to evaluate the impact of previous decisions and validate strategic outcomes. *Leading indicators*, by contrast, anticipate future trends and support proactive interventions. Lastly, *Diagnostic indicators*—which are neither purely predictive nor retrospective—provide insight into the current operational "health" of processes, enabling real-time monitoring and adjustment.

Regarless the type of classfication, KPIs must meet clearly defined criteria. One of the most commonly adopted frameworks is the SMART model, which states that indicators should be:

- **Specific:** clearly defined and focused on a particular objective or process, avoiding ambiguity.
- **Measurable:** based on quantifiable data, allowing for consistent tracking over time.
- **Attainable:** realistic and achievable within the constraints of available resources and operational context.
- **Relevant:** aligned with the organisation's strategic goals and decision-making needs.
- **Time-bound:** associated with a defined timeframe for evaluation, enabling timely assessment and action.

According to Ishak et al. (2020), approximately 70% of large private organisations report improved performance outcomes when applying this approach. However, while the SMART framework provides a solid foundation, it does not guarantee that indicators are actionable—that is, that users know what action to take when performance declines and have the means to act. To address this limitation, Parmenter (2010) proposes an alternative set of characteristics that define what makes a KPI actionable and effective, as illustrated in Table 2.1.

Table 2.1: Key attributes of effective KPIs

Attribute	**Description**
Frequent	Monitored regularly to allow timely detection of issues and prompt action.
CEO attention	Attracts the attention of senior leadership and prompts immediate follow-up when trends deviate.
Actionable	Clearly indicates what needs to be done—staff can interpret the signal and act accordingly.
Team-owned	Tied to a specific team responsible for and capable of influencing its outcome.
High impact	Influences critical success factors and affects multiple areas of business performance.
Behaviourally sound	Tested to ensure it promotes the desired behaviour and avoids unintended consequences.

Despite the solid theoretical foundation behind KPIs design, they are often used ineffectively, partly due to the lack of clear and actionable guidelines for their implementation (Franceschini et al., 2007; Parmenter, 2010).

One of the main difficulties is the growing volume of data that organisations now collect. As this volume increases, it becomes harder to choose, filter, and interpret the information that truly matters. As a result, KPIs may lose their practical value, making decision-making more influenced by how the indicators are designed, especially in terms of whether they focus on short or long-term goals. KPIs with a short term focus may lead to reactive decisions that undermine strategic direction, while KPIs focused too heavily on the long term may delay actions that require immediate attention (Franceschini et al., 2007). Another common issue is finding the right balance between having enough KPIs and not having too many. As discussed earlier, it is not about the number of indicators, but about choosing those that truly reflect strategic priorities. Nevertheless, many organisations end up with long lists of indicators that lack relevance. This can cause confusion, reduce focus, and weaken the impact of performance measurement (Parmenter, 2010).

Finally, one critical aspect that is often overlooked is the identification of Critical Success Factors (CSFs)—the key areas that must perform well for the organisation to succeed. If these are not clearly defined, or if they are confused with more general success factors, the KPIs chosen will lack direction. As Parmenter (2010) points out, “If your organisation has not completed a thorough exercise to know its CSFs, performance management cannot possibly function.”

2.5.1 KPIs for Inventory Record Inaccuracy

Within the scope of this dissertation, KPIs play a pivotal role in quantifying stock discrepancies, evaluate their impacts, and guide the implementation of effective corrective measures. Linuwih and Handayati (2025) addresses the topic of IRI highlighting two main KPIs.

Gross Variance

Gross Variance measures the total magnitude of discrepancy between recorded and actual inventory levels, without considering the direction of the error. It provides an aggregate view of inventory inaccuracy and is frequently used to assess the overall severity of the issue. This KPI helps companies benchmark inventory accuracy across sites and track the effectiveness of corrective measures.

$$\text{Gross Variance} = \sum |\text{System Inventory} - \text{Physical Inventory}| \tag{2.1}$$

Net Variance

Net Variance evaluates the overall directional bias in inventory inaccuracies. Unlike Gross Variance, this metric preserves the sign of the error and indicates whether the system tends to systematically overestimate or underestimate stock levels. A negative result may suggest a recurring underreporting issue, while a positive value may indicate systematic overstatements.

$$\text{Net Variance} = \sum (\text{System Inventory} - \text{Physical Inventory}) \tag{2.2}$$

Chapter 3

Problem Statement

This chapter will address the challenges associated with the current *SP inconsistencies* regularisation process, focusing on its structural limitations and operational inefficiencies.

Section 3.1 presents the problem context, emphasising the critical role of accurate stock management in operations and the consequences of unresolved inconsistencies within NOS. Section 3.2 identifies the main causes of stock inconsistencies in the current organisational environment and provides a detailed analysis of the existing resolution process. Section 3.3 offers a critical reflection on the inefficiencies observed in the process, while Section 3.4 outlines the key opportunities for improvement.

3.1 Problem Context

Supply chains are the backbone of today's global economy, enabling the efficient and coordinated movement of goods, information, and services across increasingly complex and interconnected networks. Their performance has a direct impact on organisational agility, cost efficiency, and customer satisfaction.

Over time, supply chains have evolved from linear, centralized structures to more dynamic and decentralized ecosystems involving multiple actors, such as manufacturers, warehouses and logistic partners. While decentralization offers greater flexibility and scalability, it also introduces challenges, particularly regarding information consistency and operational control. One of the most relevant and persistent issues in this context is Inventory Record Inaccuracy. As goods move through various points in the chain, ensuring alignment between physical inventory and recorded data becomes increasingly difficult.

This project focuses on one of the current NOS's stock control mechanisms, specifically the process for handling *SP inconsistencies*. These discrepancies can occur throughout the logistical flow, which begins at NOS's central warehouse, extends through the supply to SP warehouses, and concludes with Installation technicians delivering equipment to end customers, as illustrated in Figure 3.1. The multiplicity of stakeholders and the fragmentation of control points make this process particularly vulnerable to the occurrence of stock inconsistencies.

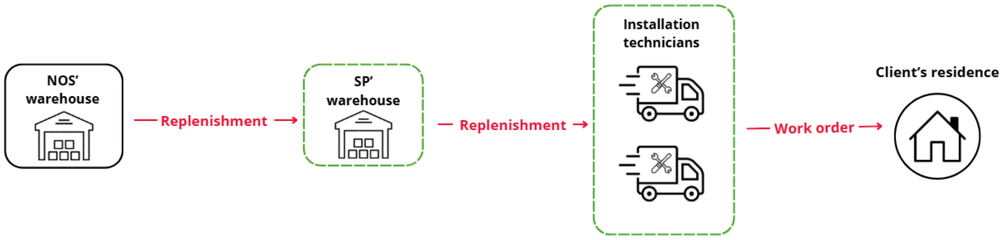

Figure 3.1: Flow where *SP inconsitencies* may arise

Despite successive automation efforts, the current stock regularisation process remains highly manual, time-consuming and fragmented. This is largely due to the need to reconcile data across multiple systems and to coordinate closely with both Technical teams and external SPs. Although these limitations are widely recognized, the process has remained virtually unchanged for over a decade, having transitioned through several management teams without being subject to a critical reassessment. This inertia can be partly explained by the fact that the process is unintuitive and operates in relative isolation from the team's core responsibilities. Additionally, it is executed and fully understood by only one team member, which limits organisational visibility and knowledge transfer.

More recently, questions have been raised regarding the actual value of this regularisation process and its contribution to overall organisational performance. In response, this study adopts a comprehensive approach, combining direct observation, process mapping, and data-driven analysis, to examine the end-to-end workflow, identify inefficiencies, and assess opportunities for targeted improvements and, if justified, a full process redesign.

3.2 AS-IS

Before analysing the current process, it is important to provide an overview of the main information systems involved, describe the standard equipment activation flow — which, when correctly followed, preserves inventory integrity — and highlight the key deviations that lead to inconsistencies.

3.2.1 Systems Involved

- **Click:** Platform used for the management and activation of equipment. Whenever a device is installed at a customer's residence, it must be registered in Click, which automatically associates the equipment with the corresponding work order (WO).

- **Siebel:** A centralised platform for managing both technical and contractual customer information. It provides access to subscribed services, active devices, and a historical record of technical interventions per customer.

- **SAP ERP:** Commonly referred to simply as SAP, this is the company's centralised ERP system, used for inventory control and for tracking both the physical location and financial status of equipment.

- **ITSM (Information Technology Service Management):** A internal platform used by NOS to manage the reception, tracking, and resolution of support tickets submitted by SPs.

3.2.2 Equipment Activation Flow and Possible Deviations

Under normal conditions, when a Service Provider receives a WO to install equipment, the activation process initiates a sequential data flow between Click, Siebel, and SAP. The technician registers the installation in Click, which triggers an update in Siebel; in turn, Siebel instructs SAP to perform the corresponding stock removal. When this sequence occurs without erros, the equipment is properly assigned to the customer's portfolio, deducted from the Partner's inventory, and no inconsistencies are generated. This process is illustrated in Figure 3.2.

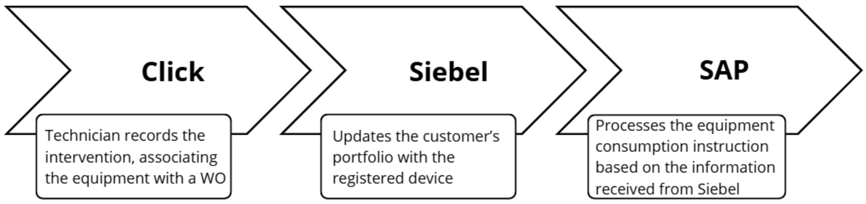

Figure 3.2: Ideal System Integration for WO registration

On the other hand, a stock inconsistency arises whenever there is an attempt to remove a device from the SP stock, but this removal is not correctly recorded in NOS's ERP system. This typically results from a failure in the activation flow. In other words, the standard chain of communication between Click, Siebel, and SAP is broken, preventing SAP from registering the stock movement.

Such inconsistencies may lead to what is known as negative inventory discrepancies, where the quantity recorded in the system is higher than the actual physical stock. When this condition persists, it results in the emergence of "ghost inventory" — stock that exists only in the system but is no longer physically available (Destro et al., 2023). It is therefore crucial to regularise these discrepancies in order to align the ERP records with the real stock situation. Accumulated unresolved inconsistencies artificially inflate the perceived inventory levels of SPs, creating the false impression that they have more equipment than they actually do. In the worst-case scenario, this leads to inefficiencies across the entire logistics flow: the mistaken belief that the SP has sufficient stock prevents replenishment orders from being placed, ultimately resulting in operational disruptions. While these scenarios illustrate the potential impact of unresolved discrepancies, it is equally important to examine their root causes. Linuwih and Handayati (2025) identify human error as a primary driver of IRI — a conclusion supported by NOS's, where recorded stock inconsistencies can be traced back to three distinct types of operational failure:

- **Unregistered installations:** The Technician physically installs the equipment at the client's residence but fails to register the WO action in Click. Without this initial registration, the entire flow is interrupted, and SAP never receives the instruction to process the stock removal.
- **Non-parametrizable swaps:** The Technician replaces a device with another from a different product family (e.g., a "V1" model is replaced with a "V2" model), violating business rules. As the system is configured to expect a specific device type, the substitution results in a mismatch that prevents proper stock removal from being recorded.
- **Partial execution of the WO:** This type of error typically arises when a WO needs to be adjusted during the installation process. For instance, although several devices may be associated with the WO, the customer may reject one of them upon the Technician's arrival, stating that it is no longer needed. In such cases, the WO must be manually updated to reflect the actual installation. If this correction is not made, Siebel blocks the order processing and fails to transmit consumption data to SAP, meaning that none of the devices, not even those successfully installed, are removed from stock.

Depending on whether the WO was registered in Click, inconsistency cases may arise through two distinct channels: the **SAP Inconsistencies Report**, when the WO is registered; or the **ITSM** platform, when it is not. Since the monthly volume of cases reported via ITSM is significantly lower than those detected through the SAP Inconsistencies Report—and given that their resolution process differs substantially—these cases were excluded from the scope of this analysis.

Figure 3.3 illustrates the disparity in case volume. It is important to note that records are only available from May onwards, as no historical data existed for the SAP Inconsistencies Report prior to that date.

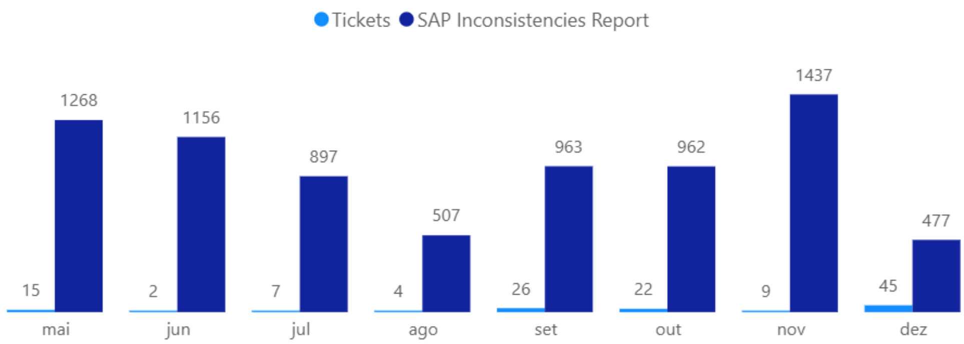

Figure 3.3: Monthly units handled via ITSM and SAP Inconsistencies Report in 2024

3.2.3 SAP Inconsistencies Report Workflow

This is the primary mechanism for identifying *SP inconsistencies*, which are detected through a parallel check between Click—where an activation attempt for the equipment is recorded—and

SAP, where no corresponding stock movement is registered. All such cases are stored in a dedicated database, which serves as the foundation for the weekly reports used in the analysis.

To promote transparency and clarify stakeholder responsibilities, the swimlane diagram in Figure 3.4 maps the current workflow for addressing these inconsistencies. This visual representation is complemented by a detailed description of the key activities performed at each stage of the process.

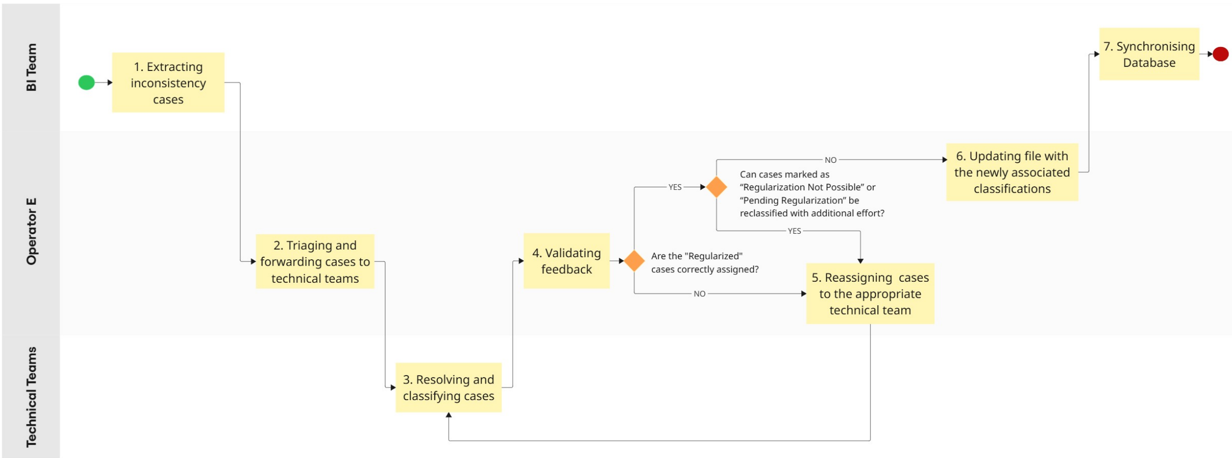

Figure 3.4: AS-IS workflow for handling inconsistencies

1. Extracting inconsistency cases

Each week, the Business Intelligence team, referred to internally as the BI team, takes a snapshot of the database containing newly identified stock inconsistencies, exports the records to an Excel file, and forwards it to Operator E. On average, the weekly file lists 407 devices, originating from the four existing SPs.

2. Triaging and forwarding cases to technical teams

On the following day, Operator E conducts an initial analysis using an automated Excel file to support this task. This preliminary analysis is not intended to result in equipment regularisation; rather, it functions as a triage mechanism, forwarding cases to three specialized technical teams. These teams, external entities outsourced by NOS, are ultimately responsible for processing the majority of regularisations. They evaluate each case individually to determine whether regularisation is feasible and, if so, proceed accordingly.

Although the majority of cases are routed externally, there are exceptions involving equipment from a specific range where regularisation is carried out directly by Operator E during this initial phase. While these cases appear to be a minority, based on observational insights, they still contribute to the overall duration of the task.

Based on an eight-week observation period, Operator E spends an average of approximately 31 minutes per execution. Of this, around 16 minutes are dedicated to running and processing the automated Excel file, while the remaining time varies depending on the volume and complexity of exceptions handled, as illustrated in Figure 3.5.

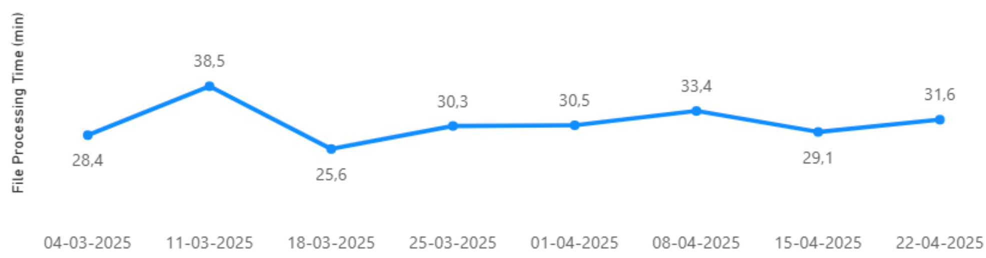

Figure 3.5: Evolution of time spent on Triaging and Forwarding cases (minutes)

3. Resolving and classifying cases

Technical teams are expected to return feedback to Operator E by the end of the week. As part of their analysis, each occurrence is assigned to one of four predefined classifications:

- **Regularised:** The equipment was regularised by the technical team, having been confirmed as correctly installed in client's residence. As a result, the stock removal it's already properly recorded in SAP.
- **No Regularisation Needed:** The equipment's historical data indicates that no corrective action is required. This often occurs when, for instance, a technician attempts to install a device, realizes it is not the supposed equipment, and replaces it before completing the operation. Although a stock removal attempt is logged, no actual physical movement occurs with the first device — therefore, no stock exit should be enforced for it. This case also highlights that not every registered removal attempt implies a genuine stock inconsistency.
- **Pending Regularisation:** Refers to cases that still require further action, such as investigation or clarification. In most situations, the technical team is unable to resolve the issue independently and ends up forwarding the case to other technical team.
- **Regularisation Not Possible:** Refers to situations where correction in SAP is not feasible due to reasons such as missing data and untraceable serial numbers.

4. Validating feedback

Once feedback is received from the technical teams, Operator E begins the feedback validation phase. This is the most time-consuming task in the process, requiring an average of six hours of effort per week. The recorded values typically fluctuate within 13 minutes of the mean over an eight-week observation period, suggesting that this stage has a relatively stable duration. The data collected are presented in Figure 3.6.

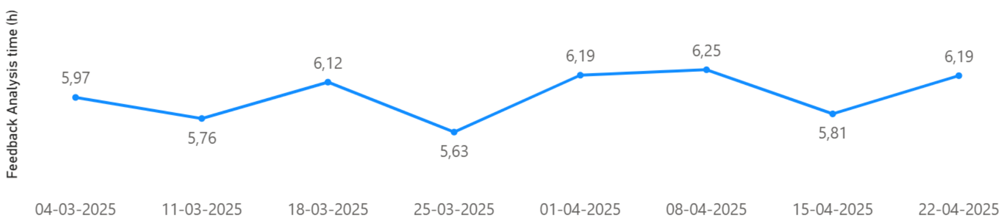

Figure 3.6: Evolution of time spent on Validating feedback (hours)

During this stage, the operator verifies two key aspects: whether the records marked as "Regularised" were correctly classified, and whether the entries labelled "Regularisation Not Possible" and "Pending regularisation" could potentially be reclassified upon further analysis. This validation step is considered critical, as it aims to reduce the number of devices remaining in these latter two categories, both of which carry distinct operational and financial implications:

- The status "Regularisation Not Possible" indicates that the equipment is deemed irrecoverable. At the semi-annual inventory close, it will be written off, marked as unusable, and the associated cost recharged to the SP. Misclassifying a device under this status can therefore result in an unwarranted financial loss for the SP.

- In contrast, the "Pending Regularisation" label suggests that NOS may not have exhausted all efforts to resolve the case. These cases can be disputed by the SP, potentially leading to credit issuance, which constitutes a financial loss for NOS.

It is also worth noting that while validation of equipment labeled "Regularised" is generally straightforward, often requiring only a quick check in SAP, cases labeled "Pending regularisation" or "Regularisation Not Possible" are significantly more complex. These require case-by-case analysis, often involving cross-referencing across multiple internal systems. Moreover, there is no standardized investigation path, as each case presents its own operational specificities.

5. Reassigning cases to the appropriate technical team

In many instances, the previous assessment leads Operator E to reassign the case, typically by sending an email—accompanied by an attached Excel file listing the cases that, although not yet resolved, are now believed to be regularizable—to either the original technical team or another team deemed better suited to resolve it. However, this often leads to follow-up communication, as the receiving teams frequently reach out to Operator E to clarify procedural or contextual details. Although informal, these interactions are time-consuming. Field observations show that phone consultations can exceed 20 minutes per case. This additional communication layer further increases the workload of what is already the most time-consuming stage in the entire workflow.

6-7. Updating file with the newly associated classifications and Synchronising the Database

Upon receiving feedback from the technical teams, Operator E updates the weekly file by modifying two key columns: one for assigning the regularisation status to each entry, and another for recording observations. The observation column includes notes from both Operator E and the technical teams, helping preserve context — particularly for unresolved cases, such as those marked as "Pending regularisation", which may require follow-up in subsequent weeks. In addition, this update enables the identification of cases classified as "No Regularisation Needed", which—once acknowledged by the BI team—should be excluded from future weekly extractions. This filtering step helps reduce rework. After all updates are completed, the file is returned to the BI team, which uses it to synchronise the database with the most recent information.

3.3 Process Diagnosis and Critical Reflection

Based on the preceding analysis, certain operational patterns raise important questions about the efficiency and added value of the current practices. The process is characterised by a significant time burden and the involvement of multiple entities operating in a non-linear flow. Frequent handovers and unproductive back-and-forth between agents result in rework. This complexity calls for a reflection on whether such a form of operation is unavoidable, or whether simplification and better integration could significantly enhance overall performance.

Empirical observation suggests that the value of the tasks described in Sections" 4. Validating feedback" and" 5. Reassigning cases to the appropriate technical team"—identified as the most time-consuming—is particularly debatable. For instance, the validation of cases classified as "Regularised" often proves redundant. Even when classifications are inaccurate, these records inevitably reappear in subsequent weekly extractions, since no actual correction is performed in the system. *—Is it therefore justifiable to invest time and resources in the manual validation of such cases?*

A similar rationale applies to the follow-up of "Pending regularisation" and "Regularisation Not Possible" cases, which typically involve attempts to determine the status of unresolved equipment and evaluate whether process acceleration or further resolution attempts are feasible through coordination with technical teams. However, considering that these teams already have the autonomy to reassign cases among themselves as appropriate, the added value of this mediation effort appears limited.*—To what extent is it justifiable to allocate resources to mediate an activity that is already decentralised and can be largely self-managed?*

At a broader level, it is essential to reflect on the proportion of total inventory affected by inconsistencies. Currently, there is no visibility into the share of stock held by partners that falls within this category, which hinders the ability to assess the actual scale of the issue. Gaining this insight is vital for evaluating the criticality of the monitoring process and understanding whether operational performance is truly at risk. Most of the questions raised call for a critical rethinking and reassessment of the process.

3.4 Improvement Opportunities

Section 3.3 offered a critical reflection on the current process. Building on that analysis, this section presents actionable improvement opportunities.

Effective improvements require clear visibility, reliable data, and a focus on root inefficiencies. With this in mind, the proposals are grouped into two areas: (i) information systems and (ii) process and operations. The first group of initiatives focuses on strengthening the digital backbone that supports the regularisation process, starting with the the implementation of a real-time monitoring tool which requires the reconstruction of the underlying data infrastructure. The second group addresses workflow design, proposing targeted interventions to reduce manual effort, simplify coordination, and improve overall performance.

3.4.1 Information System

Process Monitoring dashboard

One of the most critical shortcomings of the current process is the lack of structured monitoring and visibility. No indicators have been designed or implemented to assess either the efficiency of the process or the actual need for its occurrence. This lack of visibility prevents informed decision-making and makes it difficult to identify where and when intervention is needed. To address this, it is proposed to develop an interactive dashboard that enables real-time monitoring.

Database Restructuring

The design of any dashboard or monitoring tool must begin with the proper collection and organisation of data. A key limitation of the current setup lies in the outdated database that supports the weekly SAP Inconsistencies Reports. Developed over a decade ago, this database has become obsolete, both in terms of structure and the relevance of its stored attributes. Many fields no longer offer analytical value, which diminishes the usefulness of the information it provides.

To support meaningful data visualisation and enable more effective analysis, a new database should be developed. It must allow for historical tracking of inconsistencies, exclude obsolete fields, and apply normalization techniques to improve consistency, accuracy, and overall usability. Equally important is the automation of data population, which is essential to ensure reliability and minimize manual effort. The resulting database would serve as the informational backbone of the process, enabling continuous and effective monitoring.

Click Platform Enhancements

Improving a process also requires understanding what triggers it. Discussions with installation technicians revealed that some inconsistencies stem from the technological tools used during execution. The Click platform, used to register equipment activations, emerged as a major contributor to inventory discrepancies, making it a clear target for improvement. Two critical limitations stand out:

- **Inability to correct errors:** Click does not allow technicians to cancel or amend registrations. If a device is incorrectly logged but never physically installed, the error persists, causing stock discrepancies even if corrected on-site.

- **Lack of real-time feedback:** Registrations are asynchronous and receive no immediate confirmation from the ERP system. Errors may go unnoticed until flagged days later in the SAP Inconsistencies Report. This delay reflects a broader misalignment between physical operations and information systems, a known driver of IRI in the literature (DeHoratius and Raman, 2008).

3.4.2 Process and Operations

Workflow Optimization

The monitoring tool to be developed will enable a data-driven diagnosis of the process, supporting the evaluation of its effectiveness and the identification of improvement opportunities. In parallel, stronger integration between operational systems may help reduce the volume of inconsistencies.

However, these measures do not address a more fundamental issue: the process design itself. The current workflow is inherently time-consuming, highly dependent on coordination between multiple teams, and places a disproportionate burden on Operator E. Based on this analysis, two complementary paths may be considered. The first involves a structural reconfiguration of the process, should the data reveal fundamental flaws that justify a more substantial redesign. The second focuses on a continuous improvement approach, centred on the gradual introduction of automation and the definition of standardised rules to enhance consistency and reduce manual effort. In both cases, the objective is to streamline the process, reduce interdependencies, and ease the operational load on Operator E.

SLA Monitoring

Another improvement opportunity concerns the lack of systematic follow-up on pending cases in the weekly communications sent to technical teams. Cases classified as "Pending Regularisation" continue to appear in subsequent data extractions without any indication of how long they have remained unresolved since their initial identification.

To address this, it is essential to implement an alert mechanism that highlights overdue cases, prompting a quicker response from the responsible technical teams. At the same time, a clear Service Level Agreement (SLA) should be defined—i.e., a maximum acceptable timeframe for resolving these inconsistencies. Based on this, an operational KPI should be introduced to measure the number of cases that exceed the defined timeframe, thereby identifying SLA breaches. Making these deadlines visible within monitoring systems will facilitate the identification of delays, reinforce a sense of urgency, and support the prioritization of corrective actions by the technical teams involved.

Inventory Cycle Review

A comprehensive evaluation of the process should naturally include a critical review of the formal inventory checkpoints. While inconsistencies are addressed weekly through regularisation efforts, physical stock reconciliation — when system records are compared with actual stock — occurs only twice a year, during the May and November inventory cycles. These checkpoints remain operationally demanding. On the SP side, each cycle requires one full day of counting, totalling four days per inventory round. Although SPs perform these counts independently, Operator E plays a key role: monitoring progress, reviewing results, and investigating discrepancies. When mismatches arise — often cumulative issues from unresolved weekly efforts — the operator analyses the cases to determine whether regularisation is still possible. This follow-up frequently extends beyond the inventory day, adding significant operational burden.

The current cadence of inventory cycles may not be the most effective. Long intervals between checkpoints hinder traceability of older inconsistencies, making it harder to understand their root causes and context. Revisiting the current model — for instance, by increasing frequency or adopting a tailored approach by product category — should therefore be considered. Any changes must carefully weigh the trade-off between added effort and potential gains.

3.4.3 Summary of Improvement Opportunities

Figure 3.7 presents a summary of the main improvement opportunities identified, solution principals, and expected results.

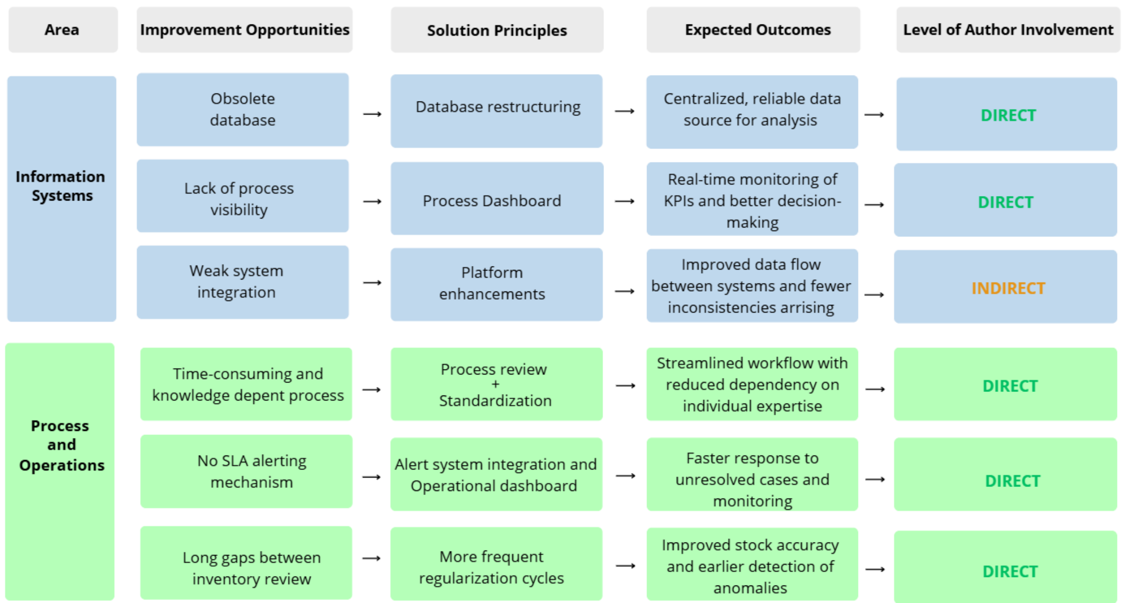

Figure 3.7: Summary of improvement opportunities

Chapter 4

Solution Development

This chapter presents the proposed solution to the challenges identified during the diagnostic phase, described in Chapter 3. The work followed two complementary development streams: the creation of a visual monitoring system and the redesign of the regularisation process to improve operational efficiency and autonomy.

Section 4.1 outlines the structuring of the database and the development of an ETL process to ensure continuous and standardised data ingestion. This phase established the foundational infrastructure for subsequent analytical work. Section 4.2 details the design and implementation of the monitoring dashboard. It covers the definition of a set of KPIs—organised into multiple analytical dimensions— model construction, and the creation of visualisations in Power BI.

Finally, Section 4.3, building on the critical assessment of the AS-IS workflow presented in Chapter 3, outlines the TO-BE regularisation process. The solution focuses on automation, redistribution of responsibilities, and integration of control mechanisms through embedded monitoring tools.

4.1 Database Restructuring

The database provided by the BI team serves as the primary data source for a future monitoring dashboard. Ensuring that this data was both reliable and well-structured was essential to support the development of meaningful analyses. For this reason, the data preparation phase was prioritised at the start of the improvement roadmap.

It was initially considered that this initiative should have been carried out in close collaboration with the BI team, given their responsibility for the weekly manipulation of the database, which also falls under their management. However, after an initial meeting held with the team to explore this possibility, it quickly became evident that such an approach was unfeasible due to three main constraints:

- **Lack of technical ownership:** The current team has limited knowledge of the database's internal structure and logic. Their involvement is mostly restricted to extracting and distributing data, as the database was originally developed several years ago by a worker who is no longer part of the team.

- **Organisational dependencies:** The database is used by several other teams across the company, each with specific needs. Changing its structure would require coordination between departments and could quickly become complex — or even unfeasible. Attributes that seem redundant to one team might be essential to another. In addition, some teams have automated processes built around the current structure, which would need to be adapted if any changes were made. These factors make it challenging to implement a single solution that suits everyone.

- **Lack of historical tracking:** The database does not retain records of devices that have already been regularised — once resolved, entries are automatically removed. However, for the purpose of building a meaningful dashboard, it was essential to maintain a consolidated historical view of all occurrences, regardless of their current status. This would allow the the team to track the time elapsed between the moment a device entered an inconsistent state and its resolution.

Given these limitations, an alternative solution was required: aggregating all weekly files stored by Operator E into a consolidated historical database. However, since these files were manually edited by multiple actors — including the technical teams and Operator E — challenges quickly emerged regarding data consistency and validation. To address these issues, an ETL pipeline was implemented to standardize inputs and generate a reliable, continuous, and maintainable data source.

4.1.1 Tools Selection

Since the dashboard was to be developed in Power BI, the initial plan was to perform the ETL using Power BI itself or, alternatively, Excel with VBA, both tools already familiar to the team.

Power BI, through Power Query and the M language, supports basic to intermediate ETL operations such as filtering, merging, and simple joins. However, the project required more advanced capabilities, including complex conditional logic, multi-step transformations, temporal tracking of records, and schema validation—features that go beyond Power BI's native functionalities. Additionally, integration with external systems like SAP is limited and often depends on specialized connectors or third-party solutions, which poses challenges for potential automated data extraction needs.

Excel with VBA, although commonly used for lightweight automation, presented several limitations as well. It offers minimal support for robust error handling, lacks scalability, becomes increasingly fragile as workflows grow in complexity, and provides no native mechanisms for ensuring data integrity. Given these constraints, both tools were ultimately deemed inadequate for building a robust ETL pipeline. As a result, a more flexible and scalable solution was implemented using Python. The Python ecosystem offers the versatility, readability, and maturity required for complex data processing. It provides powerful libraries suitable for ETL development and is compatible with advanced processing engines such as Apache Spark and AWS Glue, should future scalability needs arise (Mbata et al., 2024). Although widely used in large-scale environments, these engines were considered overly complex for this project. Given a single, well-defined data source and moderate transformation needs, a simpler and more tailored approach was adopted.

For the final stage of the ETL pipeline—storing the structured outputs—Microsoft Access was selected as a lightweight yet efficient storage solution, well-suited to the project's scale and operational requirements.

4.1.2 ETL Architecture

Before delving into the technical implementation, it is important to recognize that ETL architectures can vary significantly, depending on the objectives and the complexity of the data environment. As such, any ETL process must be carefully designed to align with the specific requirements and constraints of the project. In this case, a batch-oriented architecture was chosen, meaning that data is processed from source to destination at predefined times (Mbata et al., 2024), in this case weekly. The outline of the process can be seen in Figure 4.1, followed by an explanation.

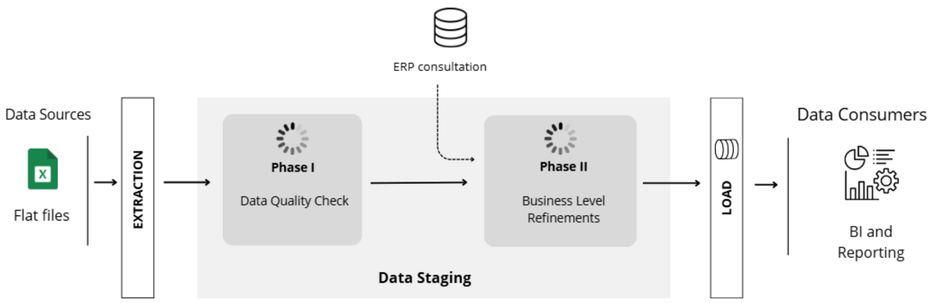

Figure 4.1: ETL architecture

Extraction

The extraction process adopts an incremental ingestion strategy, whereby only the most recent SAP Inconsistencies Report is retrieved and processed during each execution. Before proceeding, a structural validation step ensures that the file adheres to the expected schema. Any file that does not comply with the predefined column structure is excluded from further processing. Once validated, the data is loaded into a staging area, characterized as an intermediate environment where data is subsequently transformed and cleansed in successive steps (El-Sappagh et al., 2011).

Phase I: Data Validation and Quality Assurance

This phase of the pipeline aims to ensure that the data meets the required standards of quality, consistency, and structure for analytical use. Most operations at this stage were developed using the `Pandas`[1] library, a powerful tool for data manipulation and analysis. The process involves several key steps:

- **Dimensionality reduction:** Removal of attributes deemed irrelevant for the analysis. This step highlighted the extent of the original database's obsolescence, as the number of attributes was reduced from 49 to 20.

- **Type conversion and normalization:** Conversion of attributes to their appropriate data types, ensuring consistency across the dataset.

- **Structural validation:** `Pandera`[2] library was used to define a formal schema that specifies the expected data types, nullability constraints, and structural integrity of the dataset. This allowed for the early detection of issues such as missing mandatory values.

- **Data Inconsistencies Treatment:** Resolving inconsistencies in text attributes manually edited by the teams and operator E. The most common spelling and formatting errors were anticipated to ensure categorical coherence and consistency.

- **Missing value handling:** Application of multiple strategies, including constant-value imputation, conditional imputation, or rule-based logic derived from business requirements.

- **Duplicate Record Identification and Removal:** Identifying and eliminating duplicate records, avoiding redundacy in the target system.

At the end of Phase I, the data is loaded into an intermediate database, which aims to consolidate raw data with an initial processing step.

Phase II: Business-Level Aggregations and Refinements

The second phase of the pipeline focuses on reorganizing and enriching validated data, aligning it with the project's analytical goals. During this stage, all data stored in the intermediate database undergoes the necessary transformations to ensure it is properly prepared for subsequent business analysis.

Among the various actions carried out, additional attributes were created to support a more efficient and contextualized interpretation of the information. For example, one attribute was developed to enable temporal tracking of observations. This required defining a composite key capable of uniquely identifying each new record in the database. From that point on, it becomes possible to distinguish between the first occurrence of a case and its recurrences over time. Following the same logic, another column was created to flag the most recent record for each key, making it easier to identify the current state of an occurrence without relying exclusively on date intervals. These enhancements were introduced to minimize the number of calculated columns

[1] `https://pandas.pydata.org/docs/user_guide/index.html`
[2] `https://pandera.readthedocs.io/en/stable/dataframe_models.html`

and transformations required later in Power BI, thereby improving both model performance and maintainability.

Additionally, SAP integration was performed using transaction IQ09, which allows querying the current status of each device. The implementation used the `PyWin32`[3] library for SAP scripting automation and data extraction. This approach enabled correction of misclassifications assigned to certain records — often caused by human error during manual status updates or classification. For instance, some devices appear in the database with the classification "Pending regularisation" but are later confirmed in SAP to be already regularised. This verification ensures that assigned classifications accurately reflect system reality.

Loading

After the transformation phase, the process advances to data loading. The export follows a full load strategy, replacing each output file entirely on every execution, a viable approach given the manageable data size. The resulting attributes from the final processing are listed in Table A.1 of Appendix A. It is also worth noting that data is loaded denormalized, with normalization handled in Power BI.

Error Handling

To close the ETL cycle, error handling plays a crucial role in ensuring the robustness and reliability of the process. Although failures are not expected, they can occur at any stage. To address this, each time the pipeline runs, an automated email is sent to the user with a summary of the execution outcome, as illustrated in Figure A.1 in Appendix A. This feedback mechanism enables quick verification of whether the execution completed successfully or if any errors occurred that prevented it from running.

Automation of the ETL execution

To ensure regular data updates and minimise manual effort, the final version of the ETL pipeline was configured to run automatically using Windows Task Scheduler. This scheduling mechanism allows the scripts to execute on a recurring basis.

4.2 Dashboard Conception

The ETL process was specifically developed to feed the monitoring dashboard. Once the data layer was in place, a working session was held with the team to align on the dashboard's objectives and priorities. It was agreed that the tool would primarily serve a tactical function, offering middle management a clear and structured view of the process. Accordingly, the KPIs were carefully defined to reflect those priorities.

[3] `https://pypi.org/project/pywin32/`

4.2.1 KPIs Definition

Building a dashboard involves carefully defining the indicators to be analysed. A crucial task, as a poor selection can compromise the tool's overall effectiveness. Accordingly, the KPIs were grouped into three dimensions: Case Routing Efficiency, Team Performance Assessment, and Inventory Impact Assessment, each covering a different evaluation aspect. Figure 4.2 presents the main KPIs selected for each dimension, along with their corresponding descriptions. These indicators were discussed and validated with the team to ensure alignment with operational needs and analytical relevance.

It is important to note that all indicators were designed with *drill-down* and filter capabilities in mind, enabling multi-level analysis to allow users to explore data from different perspectives. Further details on the *drill-down* and filter structure are provided in Table B.1 in Appendix B.

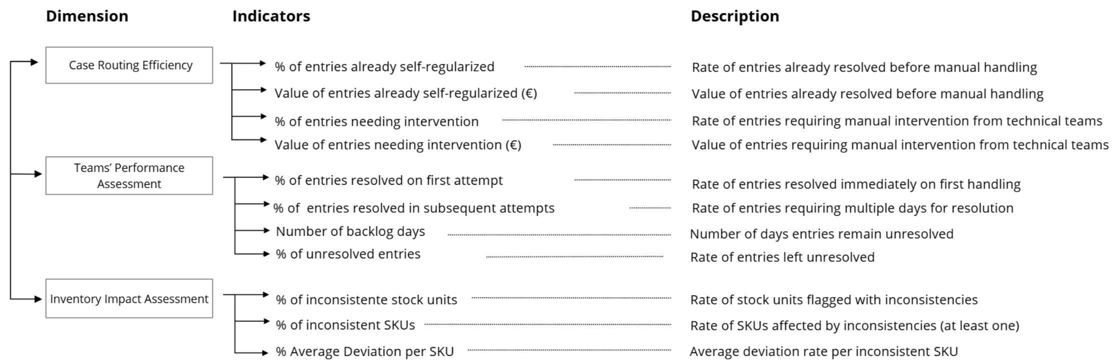

Figure 4.2: KPIs by Analytical Dimension

Case Routing Efficency

The main objective of this first dimension is to assess the relevance and timing of the current weekly initial review process. From the outset, it became important to distinguish between inconsistencies that genuinely require intervention from the technical teams and those that, at the time of analysis and routing by Operator E, are already resolved.

This issue was considered worth exploring, as direct observation revealed that the number of inconsistencies initially reported by the BI team in the weekly file was significantly higher than the number of cases actually forwarded for resolution. In many instances, inconsistencies self-corrected between the moment the data was extracted and the time it was analysed.

To reflect this reality, a set of indicators was designed to capture that difference — both in percentage terms and financial value — with the aim of supporting future decision-making. The objective is to help assess whether the current weekly frequency of analysis and regularisation should be reconsidered, based on the assumption that, over a longer interval, some cases may resolve autonomously without requiring operational effort.

Teams' Performance Assessment

The second dimension focuses on evaluating the effectiveness and responsiveness of the technical teams, responsible for resolving inconsistencies. Its primary objective is to establish a performance baseline for future comparisons—specifically between the current configuration, which, as described in Chapter 3, includes the support of Operator E (who acts as a mediator by monitoring pending cases and coordinating their resolution with technical teams), and a hypothetical scenario in which such support is no longer provided.

To this end, a set of indicators was defined to assess team performance over time, from the initial intervention to the final resolution of each case. The analysis aims to track the proportion of cases resolved on the first attempt, offering insight into the teams' ability to promptly and effectively address inconsistencies. Additionally, it monitors the percentage of cases that entered the backlog, serving as an indicator of operational inefficiency. Lastly, unresolved cases are also considered, in order to determine the technical teams' non-compliance rate—that is, cases that, despite entering the backlog, were never resolved through this weekly process.

Strategic Inventory Impact

The third dimension takes a more strategic perspective, aiming to assess the aggregate impact of inconsistencies on total stock available at the SPs. This broader view supports a more reasoned reflection on the true magnitude of the issue and whether the resources currently allocated to its resolution are proportionate. If the volume of affected stock proves limited, it is legitimate to question whether the current resolution process and the scale of effort are truly justified.

The indicators defined for this phase offer complementary perspectives, providing multiple angles of analysis and helping to avoid narrow interpretations. For example, the **% of Inconsistent Stock Units** quantifies the overall share of stock affected. This high-level metric gives a general sense of total volume impacted but does not reveal how inconsistencies are distributed across products—an important distinction, as concentrated deviations may pose different operational risks than dispersed ones.

This highlights the need for complementary indicators. The **% of Inconsistent SKUs**, for instance, assesses the spread of the issue—that is, how many distinct SKUs present deviations. A high value may indicate systemic fragility. Conversely, if the percentage is low, it becomes important to identify which specific SKUs are affected to investigate root causes.

Additionally, **% average deviation per SKU** shows the average severity of inconsistencies by indicating the proportion of each product's stock affected. This complements the previous metric by offering a more detailed view of how concentrated or critical discrepancies are within individual SKUs.

Tracking these indicators continuously allows early detection of emerging risks and supports proactive adjustments to the operational strategy.

4.2.2 Data Preparation and Model Construction

Once KPIs were established, the project advanced to the data processing and modeling phase. In addition to the data source containing historical inconsistencie, consolidated through the ETL process, two additional data sources were integrated to support the required analyses. Their integration was planned early to ensure continuous dashboard updates during development.

The first source provides the valuation of SKUs currently in stock and is extracted directly from SAP. Since this information is reviewed only once per year and follows a consistent structure, its integration into the model requires minimal effort and does not impose a significant maintenance burden.

The second source monitors the evolution of physical stock quantities across SPs over time. It was already integrated into an automated update process, given its regular use by other team members for various analytical purposes.

Once all the necessary information was gathered, the data was imported and transformed in Power Query within Power BI. Data types were adjusted, custom columns were created, and functions such as `FILTER`, `APPEND`, and `LEFT JOIN` were applied for filtering, combining, and establishing relationships between tables.

These transformations laid the groundwork for the data modeling phase, whose main goal was to normalize the dataset as much as possible, reducing redundancy, simplifying maintenance, and enabling the solution to scale over time. The data was organized into fact and dimension tables. In this case, two separate fact tables were required: one storing records of inconsistencies and another tracking available stock levels. To enable consistent analysis across both, conformed dimensions were introduced — shared reference tables aligned in structure, content, and meaning. Specifically, the "mat" table, which contains information about each SKU, and the "centro" table, which identifies the existing SPs, served as the key linking elements. Establishing these connections also played a crucial role in ensuring that filters applied in the dashboard — such as SKU selection or SP-specific views — propagated uniformly across all visualisations. Additionally, a calendar table was incorporated following best practices in Power BI modeling. It contains a continuous sequence of dates and supports reliable time-based filtering and trend analysis.

Figure 4.3 shows the data model, represented using a Unified Modeling Language (UML), developed to support the creation of the dashboard. Although the structure appears simple, several calculated columns and DAX measures were subsequently added to enable the required visual representations. The corresponding model view in Power BI is provided in Figure B.1, Appendix B.

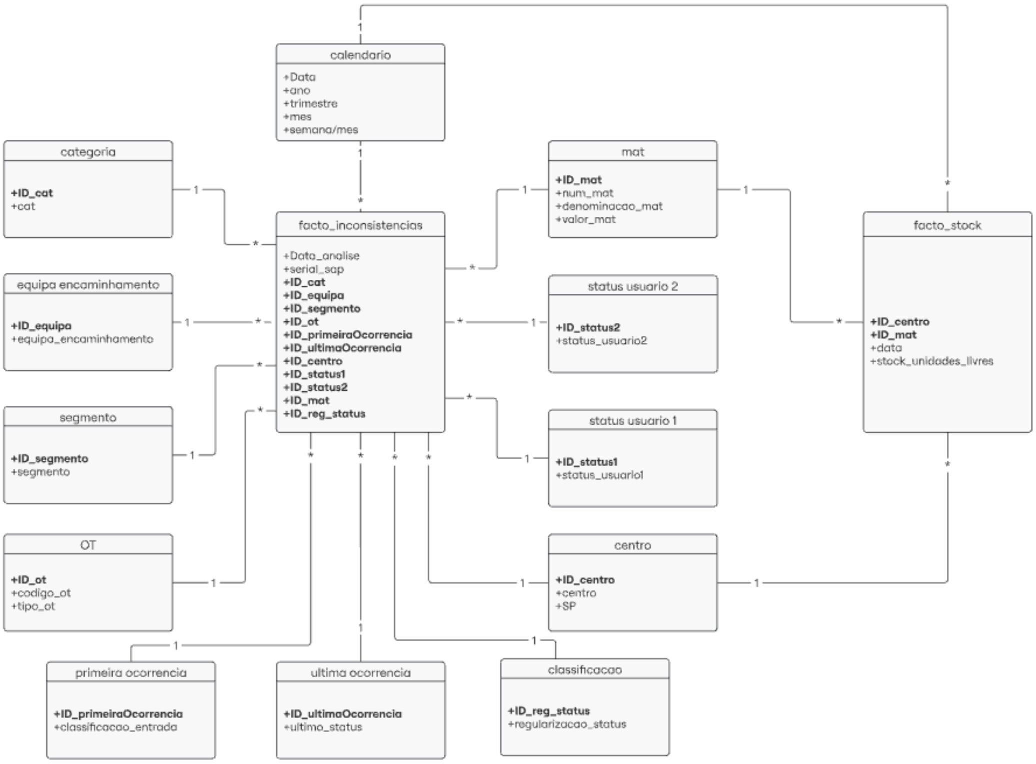

Figure 4.3: Conceptual Data Model in UML

4.2.3 Visual Layer Development

With the data infrastructure and KPIs in place, the next step was to design the dashboard layout. A visual mock-up, shown in Figure B.2 (Appendix B), was created to define the structure and organisation of the pages.

A key objective was to group the most relevant KPIs from each analytical dimension on a single page to enhance interpretability. The dashboard was thus divided into three thematic pages, each corresponding to one analytical dimension. A brief description of each is provided below, with key elements highlighted in red. For better visualisation, the developed pages can also be consulted in Appendix B.

Operational Snapshot

This first page, shown in Figure 4.4, is structured into two main analytical sections.

Section 1 provides a high-level overview of recent inconsistency entries. It highlights the distinction between cases that were self-regularised — entries which, by the time they are about to be sent to the technical teams, have already been resolved and no longer require intervention — as well as those still requiring manual handling. The information is displayed through visual cards focusing on three key aspects: case proportion, number of affected units, and financial impact, enabling a clear assessment of the operational effort involved.

Section 2 offers a snapshot of the overall process by comparing the volume of new inconsistency cases requiring manual intervention with those being closed — either through regularisation or last classification as unresolved. The analysis is presented in absolute case numbers and monetary terms, providing a clearer view of the potential operational and financial impact. This perspective supports a quick understanding of the balance between incoming workload and resolution capacity at a certain point in time. For interpretation purposes: in May 2024, a total of 1 268 inconsistencies were reported, of which 816 required manual intervention, amounting to 39 203€. In the same period, 418 cases were regularised, while 79 were closed without resolution, accounting for 3 507€.

Color coding is used strategically to enhance interpretability. For example, orange highlights cases that appear for the last time without being resolved, while blue consistently represents entries requiring manual handling. Filtering options allow users to refine the analysis by time period (year, month, or week), SP, product segment, technical team, or SKU, enabling more granular insights. For example, the SP filter helps identify the origin of inconsistencies across partners, while the technical team filter supports the analysis of case routing.

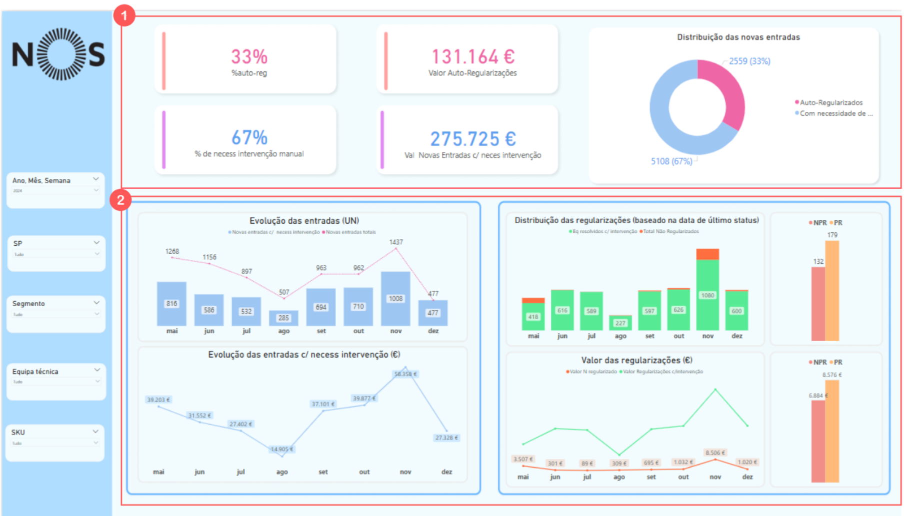

Figure 4.4: Operational Snapshot page

Teams' Performance Assessemnt

The design of this dashboard page, shown in Figure 4.5, enabled a structured and visual assessment of how technical teams handle incoming inconsistency cases. Organised into two interlinked analytical sections, the page provides a cohesive narrative of the resolution process.

Section 1 begins by contrasting the roles of Operator E and the technical teams, highlighting both the proportion and total number of entries assigned to each. It also displays the share of cases each was able to regularize on the first attempt. The analysis then narrows its focus to the performance of the technical teams, consistently represented in green. For example, a time-based area chart illustrates the monthly evolution of entries requiring intervention versus those resolved on the first attempt by technical teams, allowing users to identify trends in responsiveness over time

Section 2 maintains the focus on team performance, but shifts attention to cases not resolved on the first attempt. A distinct colour scheme highlights these cases. Within this group, the dashboard differentiates between those eventually regularised and those never resolved, using a donut chart. For the resolved cases, a bar chart segmented by time intervals illustrates the distribution of resolution delays, helping to identify patterns of backlog accumulation.

Once again, to support deeper analysis, the dashboard includes filtering options by time period, product segment, and technical team — the latter being particularly useful for performance assessments and cross-comparisons between team resolutions.

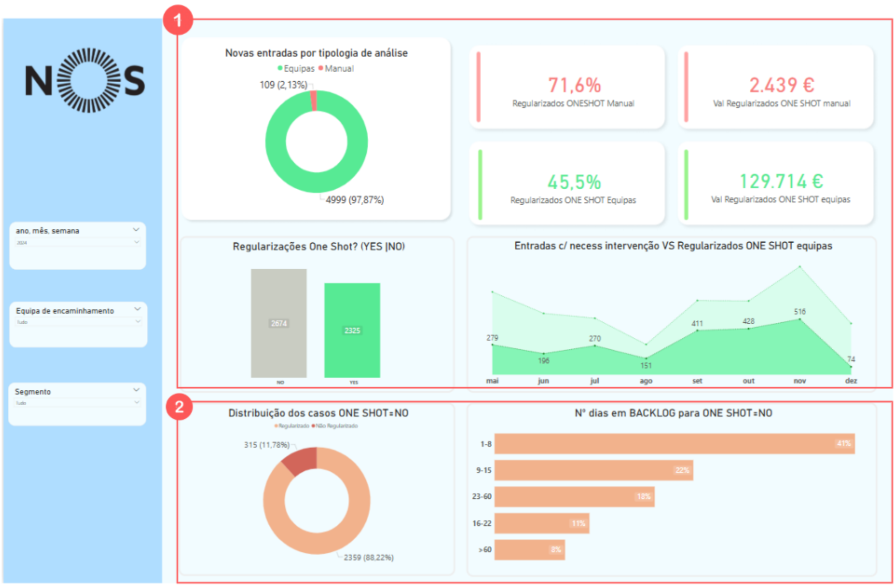

Figure 4.5: Teams' Performance Assessment Page

Inventory Impact Assessment

The final dashboard page, shown in Figure 4.6, focuses on the analysis of stock quality and how it may potentially be impacted by the emergence of new inconsistencies.

Section 1 concerns all incoming inconsistencies, regardless of their eventual resolution. The goal is to estimate their potential impact on stock at the time of detection. Section 2, by contrast, presents a view of these inconsistencies after the regularisation process has taken place. This allows for a distinction between what initially appeared to pose a risk to stock integrity and what that risk effectively translated into.

For each section, a visual card presents the % Inconsistent SKUs within the selected time period. This is followed by a temporal analysis using line charts, which display the % of Inconsistent Stock Units relative to the total stock.

Finally, a supporting table provides detailed SKU-level information, including the absolute number of affected units and the percentage of average deviation within each SKU. This visual element is intended to assist in identifying the most critical products and detecting recurring patterns that may warrant further investigation.

In this visual representation, it is important to understand how the time filters operate. Since stock data is recorded as weekly snapshots—the highest level of granularity available—selecting a broader time range (e.g., a month) will display the average of the weekly stock values within that period.

Figure 4.6: Inventory Impact Assessment Page

4.3 Process Redesign

Up to this point, the project has focused on building a solid understanding of the current *SP inconsistencies* regularisation process and developing mechanisms to support its monitoring. From this stage onwards, the investigation shifts toward identifying opportunities for process improvement.

Regardless of the specific changes to be implemented, the overarching objective is clear: to increase the autonomy of technical teams in handling cases, thereby reducing the operational burden on Operator E; to eliminate redundant and non value added activities; and to reinforce control mechanisms and their visibility. The latter partly addressed through the dashboard developed earlier.

As will be demonstrated, achieving these goals requires a profound paradigm shift that extends beyond incremental adjustments. It involves redefining roles and responsibilities, restructuring tasks, and integrating automation solutions. These transformative changes are consistent with the principles of Business Process Redesign, as outlined by Wang et al. (2024).

4.3.1 Redesign Approach Adopted

To support the implementation of the required changes, a structured approach to process redesign was followed. This approach, based on the framework proposed by Fehrer et al. (2022), breaks down the redesign into five key phases, as illustrated in Figure 4.7.

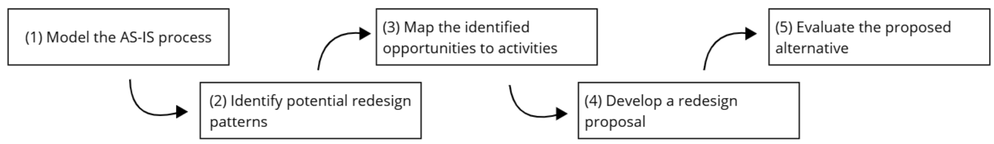

Figure 4.7: Process for applying and evaluating redesign patterns. Adapted from: Fehrer et al. (2022)

Phase 1, corresponding to the modelling of the AS-IS process, was fully developed in Chapter 3. Regarding phase 2, identifying redesign patterns, there are several approaches that are considered effective. Some of them will support this project and are described by Jansen-Vullers and Reijers (2005):

- **Task elimination**: Removing non-value-adding activities from the process in order to simplify workflows, reduce delays, and focus resources on essential task.
- **Automation**: Replacing manual tasks with digital or technological mechanisms, whenever this leads to gains in efficiency or a reduction in errors.
- **Empowerment**: Granting greater autonomy to operational actors by reducing unnecessary hierarchical approvals. The goal is to accelerate routine decisions within the teams already responsible for execution.
- **Control Relocation**: Reassigning verification or decision-making tasks to actors who are better positioned to perform them efficiently, in order to improve responsiveness and agility.

Once the main redesign patterns had been identified, their application was structured into two complementary stages: one focused on reconfiguring existing process activities, and the other on introducing new ones to address operational gaps. Each of these stages is detailed in Subsections 4.3.2 and 4.3.3.

4.3.2 Reconfiguration of Existing Activities

A review of the AS-IS process was conducted to assess how the selected redesign patterns could be used to improve the execution of existing activities. The goal was to enhance what was already in place by eliminating inefficiencies or introducing targeted optimisations. The following three activities were identified as candidates for redesign. For each, the applied pattern and the rationale behind the change are outlined.

- **Triaging and Forwarding Cases** (Control Relocation): The responsibility for weekly case assignment to technical teams was transferred from Operator E to the BI team. To operationalise this shift, a new column was added to the extraction database, applying a rule-based formula to automatically identify the appropriate technical team for each inconsistency. This change enabled the direct routing of cases to the technical teams, who can easily recognise the ones relevant to them, bypassing the need for manual triage by Operator E.

- **Validating Team Feedback** (Task Elimination): As discussed in Chapter 3, this activity was found to add limited value and appeared redundant within the overall workflow. As a result, it was removed from the process. Its elimination, however, did not compromise oversight, as monitoring continued through the tactical dashboard, which supports ongoing performance tracking. Moreover, in Chapter 5 a subsequent evaluation was conducted to assess the impact of this removal and validate the soundness of the decision.

- **Reassigning Cases to the appropriate technical team** (Control Relocation): Technical teams were given greater autonomy and were encouraged to manage case handovers directly among themselves, replacing the previous reliance on Operator E as an intermediary.

4.3.3 Integration of New Activities

The second stage focused on identifying new activities that could be added to the TO-BE process. These additions were intended to strengthen the redesigned workflow by addressing gaps and reinforcing key operational areas.

- **Sending SLA Monitoring Dashboard** (Automation) : This new activity involved sending weekly dashboard-based alerts to each technical team, offering a clear operational snapshot, particularly of pending inconsistencies and those exceeding the agreed SLA. The goal was to enhance visibility over ongoing work, allowing them to monitor progress and prioritise older cases. This approach promotes both autonomy and accountability. Further implementation details are provided in Section 4.3.5.

- **Updating SLA Monitoring Database** (Partial Automation) : To ensure the weekly delivery of SLA monitoring dashboards to the technical teams, it became necessary to update the corresponding data regularly. This activity was implemented as a partially automated task, supported by a VBA-based solution, and therefore did not introduce any significant workload for Operator E. Once again, a more detailed explanation is provided in Section 4.3.5.

Table C.1 in Appendix C provides a consolidated summary of all activities affected by the process redesign. It includes all the tasks that were restructured, eliminated, or newly introduced in the TO-BE process, detailing the operational rationale behind each change.

4.3.4 Final Redesigned Process Proposal

After incorporating all the proposed changes, the redesigned process for handling *SP inconsistencies* is illustrated in Figure 4.8. A concise description is provided below to illustrate the current practical operation of the process.

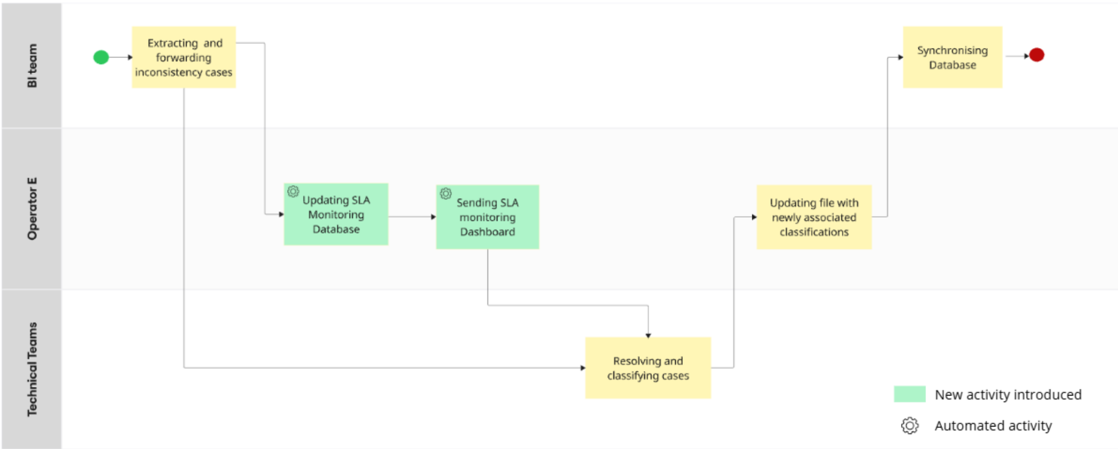

Figure 4.8: Redesigned process for managing *SP inconsistencies*

The redesigned process has become significantly more streamlined and decentralised, with Operator E now primarily responsible for tasks that are largely automated.

The workflow begins with the BI team assuming full responsibility for extracting and distributing the relevant cases to both the technical teams and Operator E.

Upon receiving the new entries, Operator E updates the database that feeds the SLA monitoring dashboards, using an automated solution based on embedded VBA logic. Once the data is refreshed, the SLA monitoring dashboards are automatically distributed to the technical teams according to a predefined schedule.

The technical teams are expected to review the dashboards before initiating the regularisation process, allowing them to prioritise their actions accordingly. During the resolution phase, these teams operate with full autonomy, including reassigning cases among themselves when necessary.

Once cases are resolved, the teams submit their feedback to Operator E, who consolidates it into the weekly tracking file. This is the only remaining manual task under his responsibility and it plays a crucial role: it ensures that structured information is preserved and passed on to the BI team for integration into the central database. This updated dataset also feeds the automated ETL pipeline that refreshes the tactical dashboard.[4] The evaluation of the proposed solution is detailed in Chapter 5, which also presents the predicted time and cost savings. Based on the positive outcomes of the pilot, the solution was deemed effective, and no further redesign iterations were required.

[4] The ETL activity is not represented in the swimlane diagram, as it belongs to a parallel monitoring stream rather than the inconsistency resolution workflow.

4.3.5 SLA Monitoring Mechanism

The introduction of the SLA monitoring dashboard mechanism in the TO-BE process was driven by the need to maintain operational control in a more automated, decentralised setup, where Operator E no longer plays a central coordination role. To address this, an operational dashboard was developed with the aim of being shared weekly with each technical team. This tool provides full visibility into the status of assigned cases, supporting both self-monitoring and greater autonomy by enabling the timely detection of delays. The dashboards designed for each team can be consulted in Appendix C.

Regarding the dashboard key features, they are described below and can be followed in Figure 4.9.

1. **Weekly overview of new entries**: This section compares the number of inconsistencies extracted in the current week with those from the previous one. It helps teams track workload evolution and detect fluctuations in the volume of inconsistencies.

2. **Distinction between new and recurring cases**: Flags whether each inconsistency is a first-time entry or a recurrence, supporting the assessment of the effectiveness of past corrective actions.

3. **Backlog time distribution (for recurring cases)**: Displays the distribution of backlog times for unresolved recurring cases, segmented into time intervals. This allows teams to detect significant delays and prioritise resolution accordingly.

4. **Identification of cases outside the SLA**: Automatically flags cases that have remained unresolved for more than 30 days. This threshold was defined as a reasonable timeframe for teams to act on assigned inconsistencies, as it typically represents four full weekly reporting cycles — and therefore four opportunities for review, resolution, or reassignment.

5. **Support table with technical details**: Displays key information for each recurring case (e.g., serial numbers and SKU codes) and supports filtering for SLA breaches, facilitating task prioritization.

Dashboard Update Process

These dashboards are powered by a secondary database, created by duplicating the raw database produced in Phase I of the ETL process. This simplified version is updated weekly in append mode, using minor transformations. An Excel template with automated VBA macros supports the update. The choice of tool was intentionally left to Operator E, who selected Excel due to familiarity and ease of use. Given the simplicity of the required transformations, this solution was considered technically sufficient. A view of the simplified interface used to update the database is shown in Figure C.1, Appendix C.

Monitoring Responsibilities

One important aspect considered during implementation was the designation of specific individuals responsible for regularly consulting the dashboards and interpreting key signals. Each technical team appointed a focal point tasked with the weekly analysis of the dashboard. These representatives were instructed to monitor key warning indicators that could signal underlying operational issues. In particular, they were advised to closely monitor the week-to-week trend in the number of recurring cases, as well as the proportion of those exceeding the 30-day SLA limit. Special attention should be given to situations where more than 50% of recurring cases breach the SLA threshold, as this may indicate structural issues requiring further investigation. Operator E was also made aware of these alert criteria and is expected to periodically verify that relevant signals are being properly addressed by the technical teams.

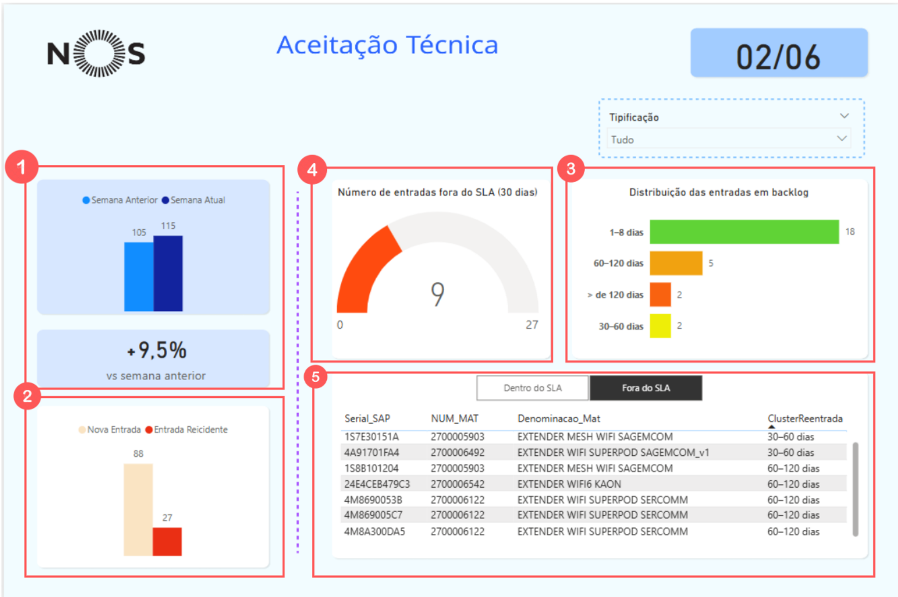

Figure 4.9: SLA monitoring Dashboard

Chapter 5

Results

This chapter outlines the key results from the implementation of the proposed improvements to the *SP inconsistencies* regularisation process and is structured across three main areas: process monitoring, workflow redesign, and the resulting time and cost savings. Section 5.1 highlights the main insights and patterns identified through the development of the tactical dashboard. Section 5.2 presents the results of the pilot project, providing evidence to validate the approaches and decisions adopted during the workflow redesign. Finally, Section 5.2.2 estimates the time and financial benefits derived from the new process.

Together, these results provide a comprehensive view of the improvements achieved and reinforce the practical relevance of the proposed solution within NOS's operational context.

5.1 Dashboard Analysis and Conclusions

Operational Snapshot Page

An analysis of the dashboard data for 2024 revealed that 33% of all new entries had already been self-regularised before any forwarding to the technical teams. When looking at weekly averages, this figure rises to approximately 37%. This finding is particularly noteworthy considering the short time window involved: there is only a one-day gap between the BI team extracting the data and the operator E performing triage and forwarding the cases. For such a narrow interval, the proportion of cases resolved autonomously is substantial, suggesting that many issues naturally resolve without requiring operational intervention. These insights prompted consideration of extending the dispatch interval to technical teams. The underlying rationale was to give more time for automatic resolution, thereby reducing the volume of cases requiring manual handling and avoiding unnecessary analysis of cases likely to self-regularise.

An attempt was made to coordinate with the BI team to test a bi-weekly dispatch. However, operational constraints quickly became apparent: since the BI team already manages weekly dispatches to multiple teams within NOS, introducing different delivery cadences would significantly increase process complexity. Consequently, the decision was made to retain the existing weekly cadence for now.

Teams' Performance Assessment Page

The second dashboard page provided a more detailed assessment of Operator E's role in the " 2. Triaging and forwarding cases to technical teams" stage. Although there was already an empirical perception that this step added limited value — and could, in fact, be efficiently executed by the BI team — the data offered objective confirmation of this assumption.

The only aspect of this stage where Operator E might have been adding value was in the direct regularisation of cases, which, as described in Chapter 3, fell within his responsibilities. However, even in this area, the impact was not very significant. In 2024, only 2.13% of all entries were assigned to Operator E, as shown in Figure 5.1, corresponding to a monthly average of around 14 cases requiring his direct intervention.

Following this analysis, a meeting was held with the team to define a revised cadence. It was agreed that Operator E would perform these regularisations on a monthly basis going forward, significantly streamlining the process without compromising case resolution.

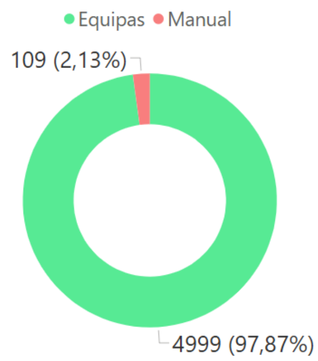

Figure 5.1: Distribution of forwarded cases between technical teams and Operator E

Inventory Impact Assessment Page

The third dashboard focused on assessing the impact of inventory inconsistencies. The analysis showed that, throughout 2024, the percentage of inconsistent stock units detected — i.e., all flagged entries relative to the overall stock — ranged between 0.12% and 0.43%, values considered to have minimal operational relevance. Furthermore, when considering only the effective inconsistencies, that is, those that remained unresolved after technical team intervention and were subsequently forwarded to the inventory cycles, the figures were negligible, with the monthly peak reaching just 0.0432%.

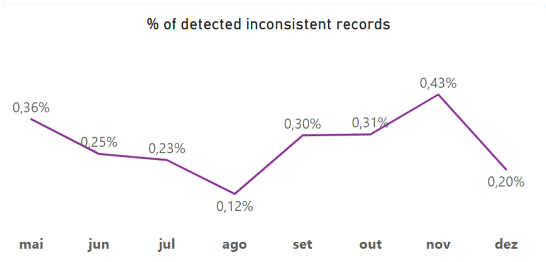

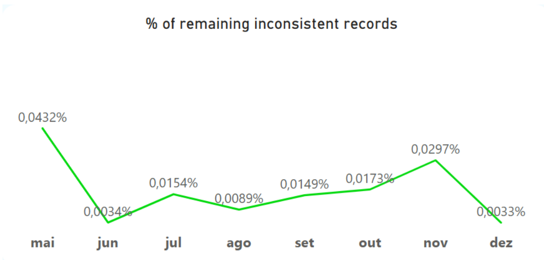

Figure 5.2: Distribution of inconsistent stock units in 2024: detected (left) vs remaining (right)

Moreover, within this dashboard page exploratory use of the available filters revealed that the SKUs with the highest % Average Deviation — i.e., those showing the greatest stock discrepancies relative to their total stock — were consistently recurring across multiple months. This pointed to a systemic pattern of inconsistencies linked specifically to a group of SKUs from the "EXTENDER" product family. Further analysis confirmed that SKUs in this category accounted for 42% of all inconsistencies reported in 2024, with a total valuation of 40 642 €.

To validate this observation, a granular monthly analysis was conducted. For each month with available historical data in 2024 (eight months in total), the five SKUs with the highest % Average Deviation per SKU were identified, forming a monthly Top 5. The recurrence of each SKU in these rankings was then assessed across the eight-month period. The results were clear: four of the five most frequently recurring SKUs in the monthly Top 5 belonged to the "EXTENDER" family. Each of these appeared in the rankings in at least seven out of the eight months analysed, as shown in Figure 5.3.

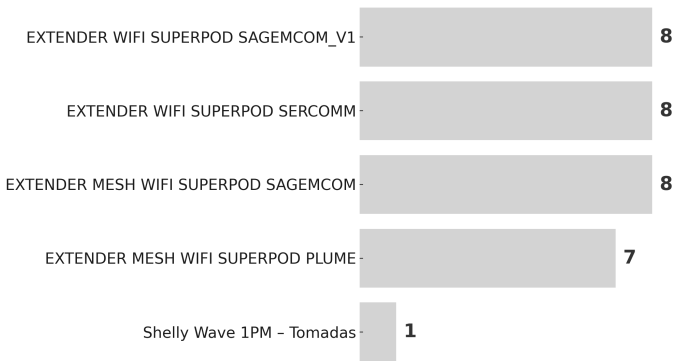

Figure 5.3: Recurrence of SKUs in the Monthly Top 5 of % Average Deviation (2024)

In light of these findings, two meetings were held: one with technical team representatives and another with SP representatives, to better understand the root cause of the issue. Following these discussions, a recurring anomaly was identified in the Click activation platform. Whenever a work order involves both an "EXTENDER" and other devices, it is common for the "EXTENDER" not to be registered, thereby generating an inconsistency entry. This issue represents yet another limitation of the platform, adding to those already identified and discussed in Chapter 3. Although the anomaly could not be technically resolved within the scope of this project, it was reported to the Supply Chain department and incorporated into awareness and mitigation sessions.

Given the recurring discrepancies associated with the "EXTENDER" product family, a second analysis was considered to evaluate the feasibility of conducting more frequent inventory checks for these SKUs. Data from 2024 showed that 43.8% of unresolved cases after successive weekly reviews were linked to "EXTENDERS", corresponding to 186 out of 315 devices—cases which can translate into operational burden during inventory reconciliation. While this reaffirmed a recurring pattern, it was not deemed sufficient to justify changes to the existing inventory strategy. As explained in Chapter 3, inventory checks are labour-intensive, requiring SP team effort and Operator E's supervision, and were therefore discouraged from being performed more frequently. Even as a targeted action, "EXTENDERS" fall within the 25–50% value range, making them less critical financially. Furthermore, financial loss for NOS only occurs when a discrepancy is formally contested by the SP and acknowledged — which, according to Operator E's feedback, is a relatively rare occurrence. For these reasons, the team chose to reinforce upstream preventive measures rather than expand inventory checks, which could undermine the efficiency gains already achieved while offering limited additional value.

5.2 Pilot Test: Assessing the Redesign Process

One of the most significant changes introduced during the process redesign was the complete removal of the task "4. Validating feedback" previously carried out by Operator E. Unlike other transformations, which involved automation or reassignment, this activity was entirely eliminated. To assess whether this had any effect on process performance, particularly on the resolution rates of the technical teams, a pilot study was conducted to compare outcomes before and after the removal of manual validation by Operator E.

A sample of 398 inconsistent cases, representing a typical weekly volume, was monitored over four weeks. Resolution rates were tracked weekly using the full sample as a baseline (398), and the results are presented in Table 5.1.

Table 5.1: Weekly resolution rates observed during the pilot

Week	Pilot Resolution Rate
Week 0	50.00%
Week 1	16.33%
Week 2	7.29%
Week 3	3.52%

To assess whether the removal of manual validation had any effect on the performance of technical teams, a non-parametric permutation test was applied. This method was chosen because it is particularly well-suited to operational settings: unlike most traditional statistical tests, it does not require assumptions of normality or equal variance, performs well with small sample sizes, and is highly flexible in handling real-world data. (Holt and Sullivan, 2023).

The hypotheses tested were defined as follows:

- **H_0 (Null Hypothesis)**: The pilot results come from the same distribution as the historical data — i.e., there is no significant effect from the change implemented in the pilot.

- **H_1 (Alternative Hypothesis)**: The pilot results do not come from the same distribution as the historical data — i.e., there is a significant effect from the change.

To build a meaningful comparison base, historical data from 2024 and 2025 (excluding May 2025, the pilot period) was reorganized into weekly entry cohorts. Each cohort grouped all cases that entered the system in a given week and was tracked over four relative weeks (week 0 to week 3), replicating the structure used in the pilot. For each cohort and week, it was calculated resolution rates. This process yielded a distribution of historical resolution rates per week, reflecting typical team performance across different time periods. The test then compared the pilot resolution rate against the corresponding historical distribution using the following logic:

- The pilot value was combined with the corresponding historical values for the same week to form a unified dataset.

- Under the null hypothesis that the pilot and historical data come from the same distribution, 10 000 random permutations were performed by shuffling all values, including the pilot, indiscriminately.

- In each permutation, one value was randomly assigned as the “simulated pilot,” with the remaining values forming the “simulated historical” group.

- The absolute difference between the simulated pilot and the mean of the simulated historical group was calculated in each iteration, generating an empirical distribution of expected differences under the null hypothesis.

- The actual observed difference, between the real pilot value and the historical mean, was then compared to this empirical distribution. The two-tailed p-value corresponds to the proportion of permutations where the simulated difference was equal to or greater than the observed one. A low p-value would indicate that such a large difference is unlikely to occur by chance, suggesting a statistically significant deviation from historical performance.

5.2.1 Pilot Test Results

Table 5.2 summarises the results of the permutation tests. The table includes: (i) the resolution rate observed in the pilot; (ii) the average historical resolution rate for the corresponding relative week (iii) the absolute difference between the two; (iv) the p-value obtained from the two-tailed permutation test.

Table 5.2: Permutation test results by week (two-tailed)

Week	Pilot	Hist. Mean	Diff. (p.p.)	p-val.
0	50.00%	25.05%	24.95	0.2067
1	16.33%	11.09%	5.24	0.7676
2	7.29%	6.40%	0.88	0.8987
3	3.52%	4.72%	1.20	0.8650

For the significance level considered (0.05), all p-values are above the threshold, meaning the null hypothesis cannot be rejected. This suggests the observed differences could reasonably be attributed to random variation rather than a true effect of the process change. In practical terms, there is no statistical evidence that removing manual validation negatively or positively impacted resolution rates across the four weeks analyzed.

5.2.2 Time and Cost Savings

To evaluate the operational and financial impact of the process redesign, an analysis was carried out focusing on changes to Operator E's weekly responsibilities, aiming to quantify the actual time and cost savings achieved. Under the new model, two manual tasks are no longer performed by Operator E: "4. Validating feedback", which required an average of 6 hours per week, and "5. Reassigning cases to the appropriate technical team", which took approximately 16 minutes per week. Although this second activity remains part of the process, it was integrated into an existing workflow managed by the BI team, already responsible for distributing cases to other NOS teams, as highlighted in Chapter 4. Thus, no additional effort or cost was incurred on the BI side, representing a clear operational gain. In contrast, a new task was introduced to maintain operational visibility: "Updating SLA Monitoring Database", to support the technical teams. This requires Operator E to update the underlying data by running an Excel file with embedded VBA, resulting in an estimated effort of 6 minutes per week. The update of the tactical dashboard, not directly related to the regularization process but a new activity, requires no additional effort, as the ETL scripts are fully automated via Windows Task Scheduler. Overall, these adjustments resulted in net weekly time savings of approximately 6 hours and 10 minutes, significantly reducing the operational burden on Operator E. In monetary terms, assuming a standard 40-hour work week, a monthly salary of 1400€, and an employer social security contribution of 23.75%, the resulting staff cost is approximately 10.00€/ h (considering an average of 4.33 weeks per month). Based on this, the estimated annual saving amounts to 3 206.40€, as shown in Table 5.3.

Table 5.3: Estimated cost savings from process redesign

	Week	Month	Year
Staff cost (€ / h)	10.00	10.00	10.00
Time saved (h)	6.17	26.72	320.64
Cost saved (€)	61.70	267.20	3206.40

Chapter 6

Conclusion and Future Work

This section presents the final conclusions and outlines opportunities for future work related to the regularisation of SP inventory inconsistencies at NOS. Section 6.1 summarises the main findings and results of the project, highlighting its operational impact. Section 6.2 outlines potential improvements, systemic limitations identified during the project, and recommendations for future initiatives.

6.1 Final Conclusions

The project was initiated with the goal of optimising the *SP inconsistencies* regularisation process at NOS. Although the topic was internally recognised as operationally relevant, there was initially a lack of clarity regarding the nature of the process itself, its operational implications, and the inefficiencies it generated. The work began with a detailed mapping of the existing workflow and exploratory interviews with key stakeholders to gain a deeper understanding of the end-to-end operation. This initial phase provided deeper insight into the executional aspects of regularisation and also revealed a broader set of improvement opportunities extending beyond the purely procedural dimension. One of these was related to the absence of a structured mechanism for process monitoring. This gap made it clear that developing a data visualisation tool would be essential for enabling continuous oversight, identifying performance trends, and supporting evidence-based management decisions. However, even before dashboard development began, a critical constraint emerged: the lack of a centralised and reliable data source. Much of the required information was fragmented and inconsistently stored. To overcome this, a local ETL pipeline was autonomously developed to consolidate data from multiple internal sources. Despite operating independently from NOS's broader data infrastructure, this pipeline proved robust and effective, enabling the project to advance with confidence toward dashboard design and implementation.

Following data consolidation, the project progressed to the definition of appropriate KPIs to support meaningful process monitoring, followed by the design of the dashboard to determine how best to represent the information visually. Several adjustments were made throughout this phase as new requirements emerged or data gaps were uncovered. Once stabilised, the dashboard significantly enhanced process visibility and enabled the extraction of key insights. Notably, its use facilitated the early identification of structural patterns in the data— particularly the recurrence of a specific subset of SKUs responsible for a disproportionate share of inconsistencies. In parallel, a temporal analysis revealed that the overall volume of stock units affected by inconsistencies was marginal. This finding helped put the issue into perspective, suggesting that the effort previously dedicated to these anomalies may have been disproportionate. It also contributed to easing internal concerns about the potential severity of the problem.

In a second phase, the project shifted its focus to workflow process redesign. During this stage, improvements were implemented based on key principles such as automation, decentralisation of responsibilities, and the elimination of non-value-added activities. Among the changes implemented, the most significant was the complete removal of the "4. Validating feedback" task, previously performed by Operator E. To assess the impact of this change, a pilot implementation was conducted over a period of four weeks. Resolution rates were monitored during this timeframe and compared to historical performance using a statistical test. The results did not reveal statistically significant differences between the pre- and post-intervention scenarios. However, this should not be interpreted as evidence that the outcomes were equivalent—only that it is not possible to affirm, that a statistical difference exists between the periods with and without the intervention role of Operator E. In this context, continuous monitoring remains critical to validate the long-term effectiveness of the intervention. Even so, the subsequent introduction of the activity "Sending SLA Monitoring Dashboard" — one of the automation mechanisms introduced — combined with increased awareness among technical teams, is expected to contribute to sustained, or even improved, resolution performance.

As a result of the process redesign, the responsibilities of the Fixed Equipment Planning Team have shifted toward higher-value tasks, primarily focused on automated monitoring and exception handling. This transition away from repetitive, manual operations has led to measurable efficiency gains. Specifically, the elimination of low-value activities resulted in a weekly time saving of 6 hours and 10 minutes, which translates into an estimated annual cost saving of 3 206.40 €.

6.2 Limitations and Future Work

Throughout the development of this project, several challenges were encountered that not only conditioned the scope of the work but also helped identify opportunities for future improvement.

One of the most immediate constraints was related to data consolidation and automation. The ETL pipeline developed played a critical role in integrating information. However, the solution currently operates locally, which imposes limitations in terms of scalability and robustness. Automation is managed through Windows Task Scheduler and requires SAP to be open during execution, creating an operational dependency that risks occasional failure. Although manual execution of the ETL remains a fallback, it entails a time burden and weakens the autonomy of the solution. In the future, migrating the pipeline to a more robust and centralised infrastructure, potentially cloud-based, would improve resilience and ensure sustainable automation.

Another key limitation stemmed from internal resistance to change. The process redesign involved eliminating tasks previously assigned to specific roles, which naturally created apprehension among some team members, particularly Operator E. This resistance delayed the pilot's launch and restricted its evaluation to a four-week period. To prevent similar barriers in future implementations, it is essential to invest in internal alignment and reinforce that optimisation efforts aim not to diminish individual contributions, but to improve overall process efficiency. Promoting openness to experimentation and establishing a feedback loop during transition phases will be critical for successful adoption—and to ensure that the organisation does not revert to outdated or inefficient practices.

In addition to the challenges faced during implementation, the project also highlighted opportunities for further investigation—particularly regarding structural limitations in the Click platform. These include asynchronous inventory updates, the inability to amend or correct work orders once submitted, and, more recently, the discovery of systematic problems affecting certain types of equipment. Addressing such limitations would require structural changes beyond the scope and resources of this project. Nevertheless, all major issues were shared with the Supply Chain department, with the intent of informing future improvements. Notably, NOS is currently developing a new platform to replace Click. It is therefore expected that the feedback collected during this project will be considered in the design of the new tool, helping to mitigate current limitations.

A further opportunity for improvement lies in reassessing the frequency at which cases are forwarded to technical teams. Due to hesitancy around changing the current cadence, alternatives were not tested. Yet preliminary analysis clearly suggested that some inconsistencies resolve naturally over time, reducing the urgency of immediate forwarding. Testing different execution intervals proved to be an interesting avenue for analysis, particularly in terms of reducing the operational burden on technical teams. Finding the right balance between timely resolution and efficient resource utilisation—without overloading the BI team—remains an area worth exploring further.

With regard to the monitoring tool developed, it could be further enhanced by integrating new KPIs that support long-term monitoring — a need that became particularly relevant upon the project's completion. This underscores the importance of periodically revisiting performance indicators to ensure they remain aligned with evolving operational objectives. One such KPI is the average resolution time, which stands out for its objectivity and its ability to track whether the process redesign and automation mechanisms have helped maintain or accelerate case handling over time. Another valuable addition would be the proportion of cases resolved within SLA. Although the current dashboard already includes team-level performance metrics—such as first-attempt resolution rate and backlog days—this SLA indicator would introduce a more strategic dimension. It is worth noting that a similar SLA metric exists in the operational dashboards shared with technical teams; however, it reflects a short-term, execution-level view that is fragmented across technical teams. Incorporating SLA compliance at the tactical level would provide management with broader visibility into overall service performance.

Bibliography

Alstyne, M. W. V., Parker, G. G., and Choudary, S. P. (2016). Pipelines, platforms, and the new rules of strategy. *Harvard Business Review*, 94(4):54–62.

Anand, A., Wamba, S. F., and Gnanouzou, D. (2013). A literature review on business process management, business process reengineering, and business process innovation. In Aalst, W. V. D., Mylopoulos, J., Rosemann, M., Shaw, M. J., and Szyperski, C., editors, *Enterprise and Organizational Modeling and Simulation*, volume 153, pages 1–23. Springer Berlin Heidelberg, Berlin, Heidelberg.

Anton, E., Duong, T., Aptyka, M., and Teuteberg, F. (2023). Beyond digital data and information technology: Conceptualizing data-driven culture. *Pacific Asia Journal of the Association for Information Systems*, 15(3):1–36. Accessed: 2025-02-15.

Antunes, P. and Mourão, H. (2011). Resilient business process management: Framework and services. *Expert Systems with Applications*, 38(2):1241–1254.

Bach, B., Freeman, E., Abdul-Rahman, A., Turkay, C., Khan, S., Fan, Y., and Chen, M. (2022). Dashboard design patterns. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–11.

Badawy, M., Ahmed, A. E.-A., Idrees, A., Hefny, H., and Hossam, S. (2016). A survey on exploring key performance indicators. *Future Computing and Informatics Journal*, 1.

Brynjolfsson, E., Hitt, L. M., and Kim, H. H. (2011). Strength in numbers: How does data-driven decisionmaking affect firm performance? *SSRN Electronic Journal*.

Costa, I. (2024). Business intelligence: Strategic advantage. KAIZEN™ article. Available at: kaizen.com.

Costa, S. and Santos, M. Y. (2012). Sistema de business intelligence no suporte à gestão estratégica: Caso prático no comércio de equipamentos eletrónicos. *Não especificado*. Centro de Investigação ALGORITMI, Universidade do Minho, Portugal.

DeHoratius, N. and Raman, A. (2008). Inventory record inaccuracy: An empirical analysis. *Management Science*, 54(4):627–641.

Destro, I. R., Staudt, F. H., Somensi, K., and Taboada, C. (2023). The impacts of inventory record inaccuracy and cycle counting on distribution center performance. *Production*, 33:e20220077.

Dijkman, R., Vanderfeesten, I., and Reijers, H. A. (2011). The road to a business process architecture: An overview of approaches and their use. BETA Publicatie: Working Papers 350, Technische Universiteit Eindhoven. Publisher's PDF, Version of Record.

Doe, P. (2022). Arena simulation of spare parts inventory management: A case of a telecom managed service provider. *EuroMed Journal of Business*, 14:2022.

Eckerson, W. W. (2011). *Performance Dashboards: Measuring, Monitoring, and Managing Your Business*. J. Wiley & Sons, Hoboken, N.J, 2nd ed edition.

El-Sappagh, S. H. A., Ahmed, A. M., and El Bastawissy, A. H. (2011). A proposed model for data warehouse etl processes. *Journal of King Saud University - Computer and Information Sciences*, 23(2):91–104.

Eyeregba, M. E., Mokogwu, C., Ewim, S. E., and Olorunyomi, T. D. (2024). Streamlining business processes through advanced process mapping tools: Applications in finance and beyond. *International Journal of Scientific Research Updates*, 8(2):137–148.

Fehrer, T., Fischer, D. A., Leemans, S. J. J., Röglinger, M., and Wynn, M. T. (2022). An assisted approach to business process redesign. *Decision Support Systems*, 156:113749.

Fernandes, F., Correia, J., and Pontes, A. (2023). Business intelligence: Trends and the impact of emerging technologies. In *2023 18th Iberian Conference on Information Systems and Technologies (CISTI)*, pages 1–6, Aveiro, Portugal. IEEE.

Few, S. (2006). *Information Dashboard Design: The Effective Visual Communication of Data*. O'Reilly Media, Sebastopol, CA.

Foley, (2010). What is business intelligence? *International Journal of Business Intelligence Research*, 1(4):1–28.

Franceschini, F., Galetto, M., and Maisano, D. (2007). *Management by Measurement: Designing Key Indicators and Performance Measurement Systems*.

Hamed, I. and Ghozzi, F. (2015). A knowledge-based approach for quality-aware etl process. In *2015 6th International Conference on Information Systems and Economic Intelligence (SIIE)*, pages 104–112, Hammamet, Tunisia. IEEE.

Hammer, M. and Champy, J. (1993). Reengineering the corporation: A manifesto for business revolution. *Business Horizons*, 36(5):90–91.

Harmon, P. (2010). *Business Process Change: A Guide for Business Managers and BPM and Six Sigma Professionals*. Elsevier, Amsterdam.

Herschel, R. T. and Clements, N. (2017). The importance of storytelling in business intelligence. *International Journal of Business Intelligence Research*, 8(1):26–39.

Holt, C. A. and Sullivan, S. P. (2023). Permutation tests for experimental data. *Experimental Economics*, 26(4):775–812.

Ishak, Z., Fong, S. L., and Shin, S. C. (2020). Smart kpi management system framework. *Proceedings of the International Conference on Information Science and Computing*.

Jansen-Vullers, M. and Reijers, H. (2005). Business process redesign in healthcare: Towards a structured approach. *INFOR: Information Systems and Operational Research*, 43(4):321–339.

Kantor, I., Bütün, H., Robineau, J.-L., Maqbool, S., Zwaenepoel, B., Norbert, R., Arias, S., Cervo, H., Wolf, F., and Santecchia, A. (2019). Long- and short-lists of key performance indicators in industry and for industrial symbiosis. Technical report, EPFL, UGent, ArcelorMittal, CEMEX, INEOS, Omya UK Ltd. EPOS Project Deliverable.

Kitchin, R., Lauriault, T. P., and McArdle, G. (2015). Knowing and governing cities through urban indicators, city benchmarking and real-time dashboards. *Regional Studies, Regional Science*, 2(1):6–28.

Linuwih, H. W. and Handayati, Y. (2025). Quantitative analysis of inventory record inaccuracy (iri): A case study on warehouse stock discrepancies. *International Journal of Current Science Research and Review*, 8(1):249–253.

Lousa, A., Pedrosa, I., and Bernardino, J. (2019). Avaliação e análise de ferramentas business intelligence para visualização de dados: Evaluation and analysis of business intelligence data visualization tools. *Nome da Revista*.

Luhn, H. P. (1958). A business intelligence system. *IBM Journal of Research and Development*, 2(4):314–319.

Lyon, W. (2019). Microsoft power bi desktop: A free and user-friendly software programme for data visualisations in the social sciences. *Historia*, 64(1).

Mbata, A., Sripada, Y., and Zhong, M. (2024). A survey of pipeline tools for data engineering. Preprint, accessed on 2025-05-12.

Melin, U. and Axelsson, K. (2016). Action in action research: Elaborating the concepts of action, roles and dilemmas in a public e-service development project. *Journal of Systems and Information Technology*, 18(2):118–147.

Microsoft (2024). Microsoft named a leader in the 2024 gartner magic quadrant for analytics and business intelligence platforms. Available at: powerbi.microsoft.com. Accessed on: May 25, 2025.

Miguéis, V. (2024). Data science for business. PowerPoint presentation, course “Data Science for Business”. Master’s in Industrial Engineering and Management.

Nkwake, A. N. (2023). Validity in performance measurement. In *Validity in Assessing Learning*, pages 131–145. Springer.

O’Brien, R. (1998). An overview of the methodological approach of action research. Unpublished paper to Professor Joan Cherry, Faculty of Information Studies, University of Toronto.

Ogunsola, K. O., Balogun, E. D., and Ogunmokun, A. S. (2022). Developing an automated ETL pipeline model for enhanced data quality and governance in analytics. *International Journal of Information Systems and Computer Sciences*, 3(1):791–796. January–February Issue.

Parmenter, D. (2010). *Key Performance Indicators: Developing, Implementing, and Using Winning KPIs*. John Wiley & Sons, Hoboken, NJ, 2 edition.

Rahman, A. A., Adamu, Y. B., and Harun, P. (2017). Review on dashboard application from managerial perspective. In *2017 International Conference on Research and Innovation in Information Systems (ICRIIS)*, pages 1–5, Langkawi, Malaysia. IEEE.

Rausch, P., Sheta, A. F., and Ayesh, A., editors (2013). *Business Intelligence and Performance Management: Theory, Systems and Industrial Applications*. Advanced Information and Knowledge Processing. Springer London, London.

Rekik, Y., Syntetos, A. A., and Glock, C. H. (2019). Inventory inaccuracy in retailing: Does it matter? Technical report, ECR Community Shrinkage and On-shelf Availability Group. Report supported by the ECR Community, Shrinkage On-shelf Availability Group.

Santo Andrade, S. H. M. (2020). Data analytics to increase performance in the human resources area. *Journal of Autonomous Intelligence*, 2(1):15–20.

Sheppard, G. M. and Brown, K. A. (1993). Predicting inventory record-keeping errors with discriminant analysis: A field experiment. *International Journal of Production Economics*, 32:39–51.

Solanki, A., Jain, K., and Jadiga, S. (2024). Building a data-driven culture: Empowering organizations with business intelligence. *International Journal of Computer Trends and Technology*, 72(2):46–55.

Wang, C.-N., Vo, T. T. B. C., Hsu, H.-P., Chung, Y.-C., Nguyen, N. T., and Nhieu, N.-L. (2024). Improving processing efficiency through workflow process reengineering, simulation and value stream mapping: a case study of business process reengineering. *Business Process Management Journal*, 30(7):2482–2515.

Zhao, Y. (2021). Transformation of business analytics from business intelligence. *E3S Web of Conferences*, 253:03013.

Zheng, J. G. (2018). Data visualization for business intelligence. In Munoz, J. M., editor, *Global Business Intelligence*. Taylor & Francis.

Appendix A

ETL Output and Execution Logs

Table A.1: Attribute dictionary of the SAP Inconsistencies dataset

Attribute	**Description**
serial_sap	Traceable serial number in SAP
num_mat	Material reference code (SKU) used to identify and distinguish product variants
denominacao_mat	Material description
centro	SP identification code to which device is shipped
deposito	Equipment's location (direct or reverse logistics)
cat	Identifies whether the equipment belongs to the fixed or mobile category
codigo_ot	Code associated with the WO
tipo_ot	Type of WO (installation vs. maintenance)
data_primeiro_status	Timestamp when the equipment first entered an inconsistency state
categoria_equipamento	Distinguishes equipment reserved for residential or enterprise use
data_status	Date when the cases were extracted from the database
regularizacao_status	Status equipment classification (e.g., Regularized)
regularizacao_observacoes	Optional user notes related to equipment regularization
sp	Identifies SP name
data_analise	Date the file was handled by Operator E
status_usuario1	Equipment status in SAP at the time of the analysis
status_usuario2	Equipment status in SAP at the time of ETL processing
classificacao_entrada	Differentiates between new and recurring entries
ultimo_status	Indicates whether this is the last appearance of the case in the database
equipa_encaminhamento	Technical team assigned to resolve the inconsistency
origem	Indicates the origin file of the SAP Inconsistencies Report
equipamento_adicional	Indicates whether the WO included one or more associated devices

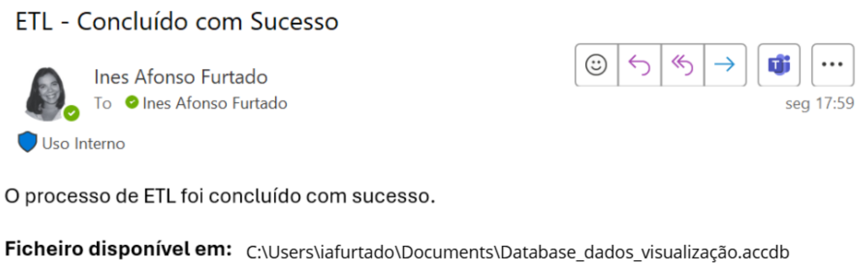

Figure A.1: ETL Execution Confirmation Email

Appendix B

KPIs and Dashboard Conception

Table B.1: KPIs Summary Table

				Interactions (drill-down and filters)				
Level	Topic	KPI name	KPI type	Time	SP	Product segment	Technical team	SKU
Process Level	Case Routing Efficency	% entries already self-regularized	Lagging	x	x	x	x	x
		Value of entries already self-regularized (€)		x	x	x	x	x
		% entries needing intervention		x	x	x	x	x
		Value of entries needing intervention (€)		x	x	x	x	x
	Teams' Performance Assessment	% of entries resolved on first attempt		x		x	x	
		% of entries resolved in subsequent attempts		x		x	x	
		Number of backlog days		x		x	x	
		% of unresolved entries		x		x	x	
Tactical Level	Inventory Impact Assessment	% of inconsistent stock units		x	x			
		% of inconsistent SKUs		x	x			
		% Average Deviation per SKU		x	x			x

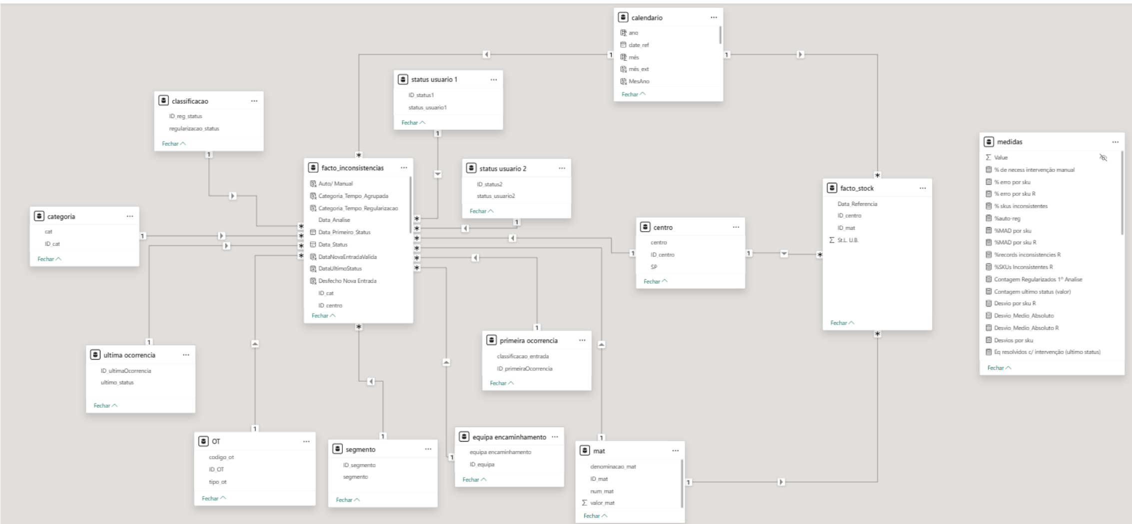

Figure B.1: Relationship Between Tables and Measures Created

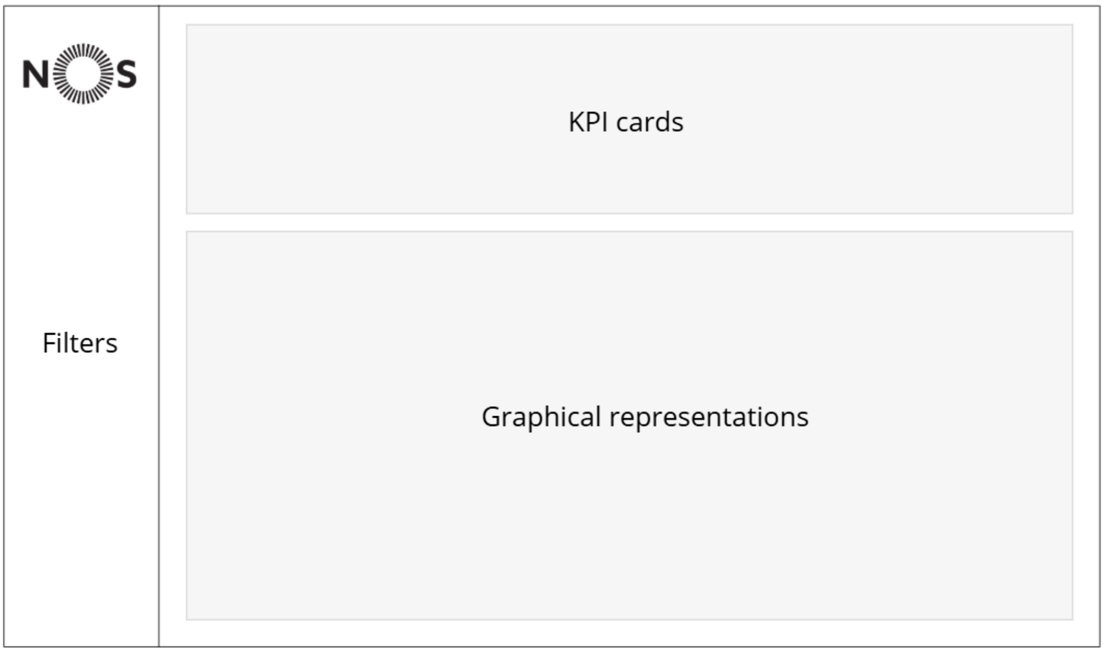

Figure B.2: Dashboard Layout

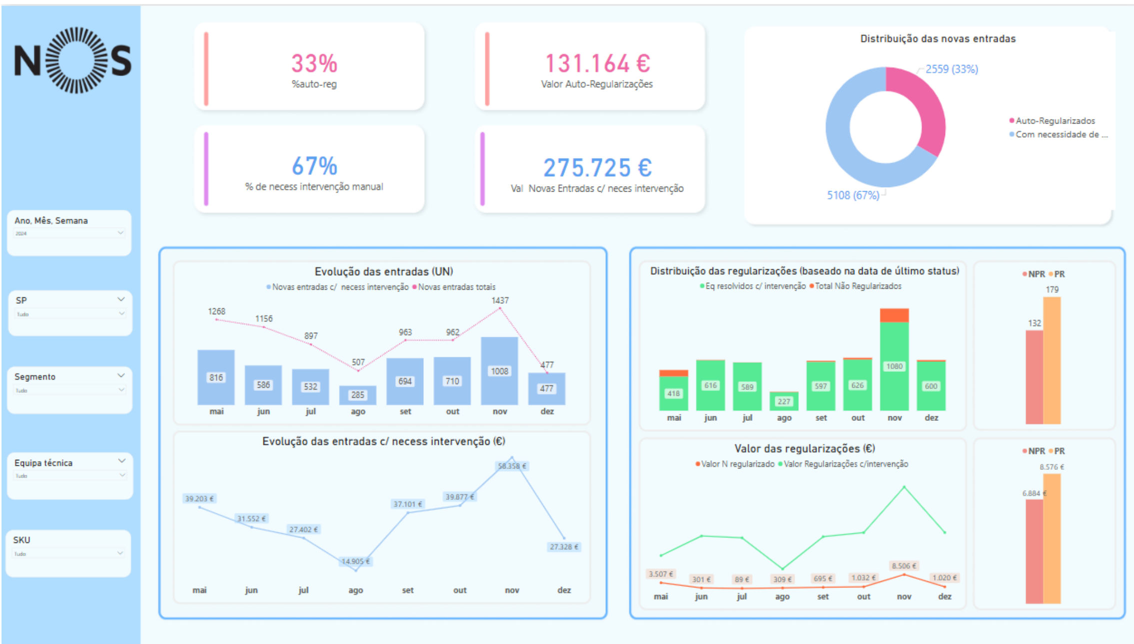

Figure B.3: Operation Snapshot Page

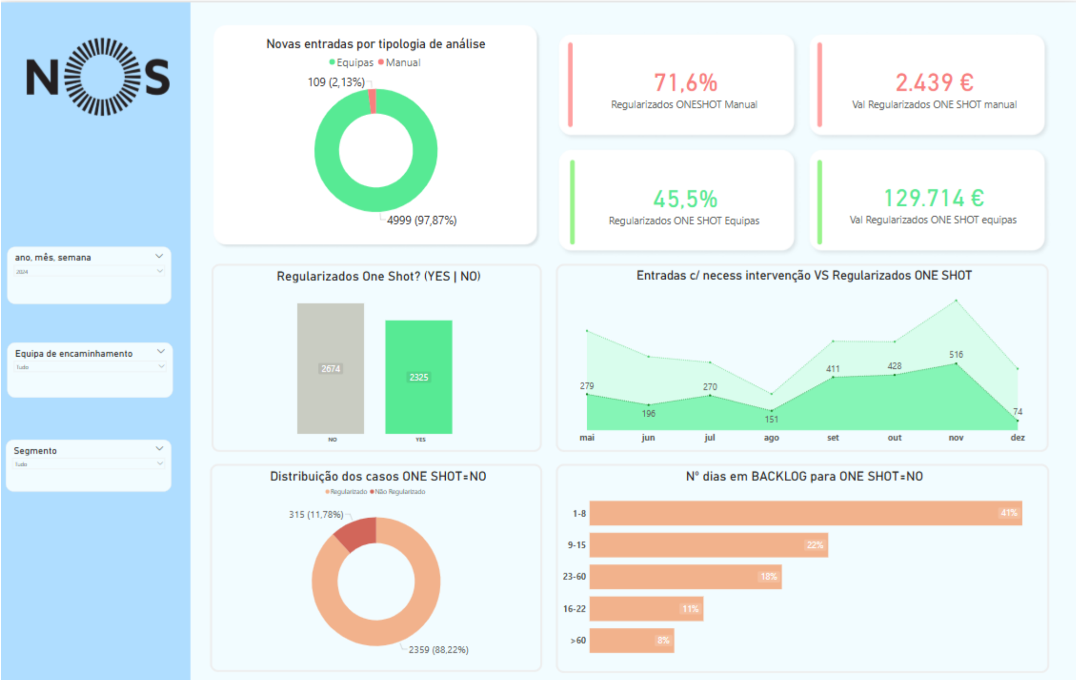

Figure B.4: Teams' Performance Assessment Page

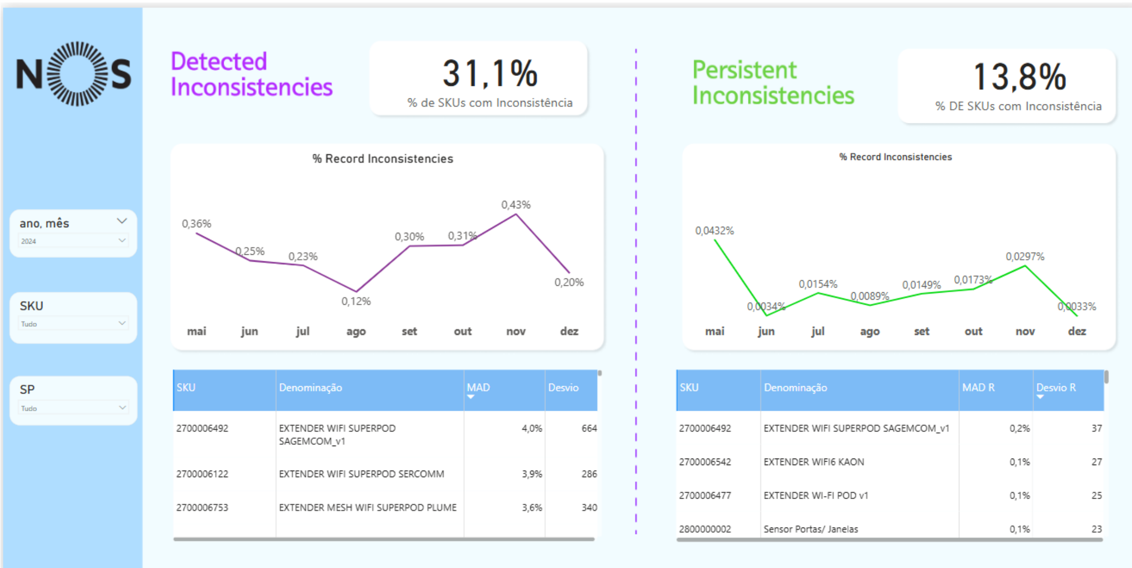

Figure B.5: Inventory Impact Assessment Page

Appendix C

Process Redesign and Associated Procedures

Table C.1: Summary of Process Changes and Applied Redesign Patterns

Activity Description	Process		Redesign Pattern	Obs:
	AS-IS	TO-BE		
Triaging and Forwarding Cases	x		Control Relocation	BI team assumes direct responsibility for case routing, reducing Operator E's workload.
Validating Team Feedback	x		Task Elimination	Removed due to low added value. Monitoring continues via dashboards.
Reassigning cases to the appropriate technical team	x		Control Relocation	Teams manage handovers internally, removing the need for Operator E as intermediary.
Sending SLA monitoring Dashboard		x	Automation	Introduction of automatic SLA alerts to enhance team visibility over case resolution.
Updating SLA monitoring Database		x	Partial Automation	VBA automation updates the data powering the SLA monitoring dashboards, reducing manual workload.

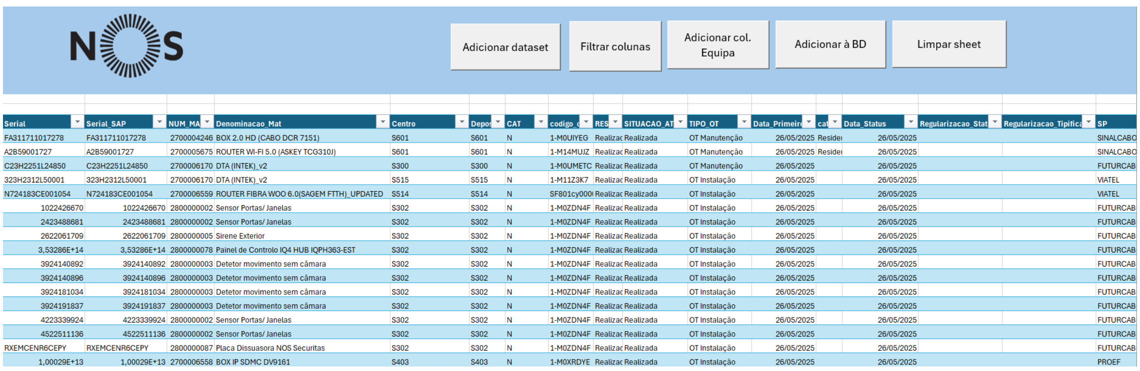

Figure C.1: Interface for Updating the SLA Monitoring Dashboard

Figure C.2: SLA Monitoring Dashboard for *Aceitação Técnica* technical team

Figure C.3: SLA Monitoring Dashboard for *Apoio Técnico* technical team

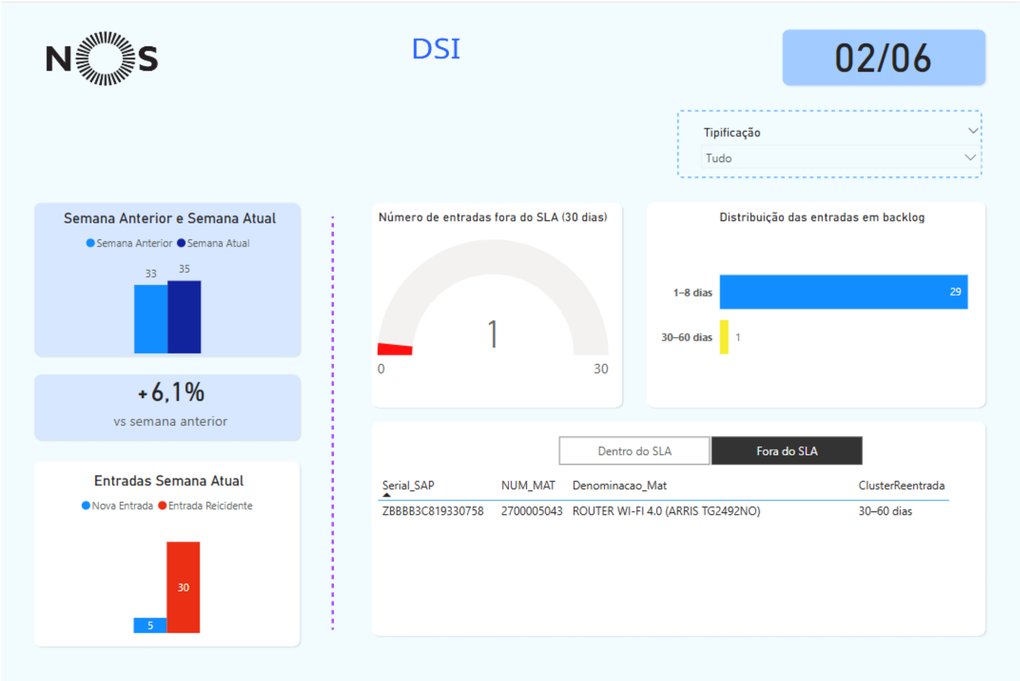

Figure C.4: SLA Monitoring Dashboard for DSI technical team

Appendix D

Project's Impact on Sustainable Development Goals

Table D.1: Project's impact on selected Sustainable Development Goals

Goal	**Target**	**Contribution**	**Indicators**
4	4.4	Development of employees' technical skills through the use of tools such as Power BI.	Greater contact with digital systems, promoting skill acquisition and analytical capacity.
8	8.2	Use of automation mechanisms, elimination of repetitive tasks, and reassignment of responsibilities, streamlining the workflow.	Higher efficiency processes.
8	8.3	Reduction of monotonous tasks, enabling staff to focus on higher-value activities.	Workload reduction and increased productivity.
9	9.5	Implementation of a data infrastructure, fostering digital transformation and improved monitoring.	Enhanced digital capabilities and agility in managing data.