Digital Integration of EHS-Compliance and Risk Management in Industrial Enterprises: The EHS-Digital Loop Authoring Platform

Tanvi Bhushan
Dhaval Sikligar

Digital Integration of EHS-Compliance and Risk Management in Industrial Enterprises: the EHS-Digital Loop Authoring Platform

Tanvi Bhushan, M.S.
Dhaval Sikligar, M.S.

Publication Info

THE AMERICAN JOURNAL OF ENGINEERING AND TECHNOLOGY (ISSN: 2689-0984)

ISBN: - 978-1-957653-55-6

CROSSREF DOI: - https://doi.org/10.37547/tajet/book-01

PUBLISHED DATE: - 02 August 2025

Preface

The accelerating convergence of digitalisation, sustainability demands, and heightened societal expectations is transforming the way industrial enterprises manage Environmental, Health & Safety (EHS) risks. Once regarded chiefly as a cost of compliance, EHS has now become a strategic lever that shapes corporate resilience, investor confidence, and licence to operate. At the same time, rapid advances in industrial IoT, cloud analytics, and artificial intelligence are unlocking unprecedented volumes of operational data yet posing new integration challenges for safety professionals, risk managers, and board-level decision-makers alike.

This monograph—Digital Integration of EHS-Compliance and Risk Management in Industrial Enterprises: The EHS-Digital Loop Authoring Platform—was conceived to address a persistent gap between traditional, paper-centric EHS frameworks and the dynamic, data-driven realities of modern production systems. Drawing upon international standards (ISO 45001, ISO 14001, ISO 31000), North-American regulatory practice (OSHA, EPA), and peer-reviewed research published over the past five years, the work proposes a coherent architecture that couples leading-indicator analytics with real-time operational decision-support. The resulting "Plan → Sense → Predict → Act" loop is positioned not as a theoretical abstraction but as a practical blueprint capable of scaling across multiple industries—from food processing and chemicals to construction megaprojects.

The text pursues three complementary aims. First, it synthesises fragmented literature on risk-based thinking, cyber-physical safety systems, and ESG performance into a single, methodologically rigorous narrative

accessible to scholars and practitioners. Second, it offers an original, formally specified meta-model that aligns the clause structures of ISO 45001, ISO 14001 and ISO 31000, thereby enabling automated governance checks and audit-ready traceability. Third, it demonstrates—through secondary case scenarios, Monte-Carlo simulation, and sensitivity analyses—how the proposed digital loop can generate economic returns, reduce carbon footprints, and enhance organisational resilience without resorting to unverified field trials.

Although the modelling results are encouraging, the author is mindful of the inherent limitations of a desk-based approach. Accordingly, each chapter strives for transparency: data sources are fully cited, assumptions are explicitly stated, and potential sources of bias are acknowledged. The reader is invited to treat the presented figures not as deterministic forecasts but as boundary-tested estimates that can be refined through longitudinal pilot studies.

This work should be of interest to EHS managers, digital transformation leaders, regulatory analysts, and researchers exploring the intersection of safety science, industrial analytics, and sustainable operations. By grounding its recommendations in publicly verifiable evidence and open modelling techniques, the monograph aspires to serve as both a scholarly contribution and a pragmatic guide for organisations embarking on their own journey towards integrated, digitally enabled EHS excellence.

Tanvi Bhushan, M.S., Dhaval Sikligar, M.S.

Tanvi Bhushan holds a Master of Science in Environmental Engineering from Lamar University, Texas, and currently serves as an Environmental Health & Safety Consultant at AARC Environmental Inc. Her expertise spans regulatory compliance, industrial risk assessment, and the digital transformation of EHS systems. She has worked across diverse sectors—construction, logistics, chemical manufacturing, and aerospace—designing compliance strategies aligned with EPA, OSHA, and ISO frameworks. Notably, she led the award-winning "Mission Zero" initiative at FreshDirect and achieved a 75% reduction in incidents at Dannick Inc.

Her key federal projects while in Dannick Inc:

1. INSTALLATION OF CHILLER PLANT, DAVIS BARRACKS
2. GRANT HALL BARRACKS RENOVATION AND REPAIR
3. CYBER ENGINEERING AND ACADEMIC CENTER
4. BRADLEY BARRACKS RENOVATION & MODERNIZATION
5. NATIONAL PURPLE HEART HALL OF HONOR IN ORANGE COUNTY NY

Her key scientific publications include:

1. "Programs for Achieving Zero Injuries and Optimizing the Infrastructure of Production Facilities" – International Journal of All Research Education & Scientific Methods (IJARESM) ISSN: 2455-6211, Volume 13, Issue 2, February-2025
2. "A Systematic Approach to Reducing Industrial Accidents and Enhancing Workplace Safety" – International Journal of Scientific Engineering and Science (IJSES) ISSN (Online): 2456-7361, Volume 9, Issue 1, pp. 77-80, 2025
3. "Enhancing System Resilience Through Leading and Lagging Indicators in Risk Management" - International Journal of Science and Research (IJSR) ISSN: 2319-7064, Volume 14 Issue 3, March 2025

Tanvi was honored with the Catalyx Ventures Award (Sustainable Champion) in 2025 for her leadership in environmental innovation. She is an active member of NEHA (National Environmental Health Association) and the American Academy of Environmental Engineers and Scientists (AAEES). Her research and fieldwork continue to bridge policy, science, and technology for a more sustainable and safer industrial future.

Contact: bhushantanvi24@gmail.com

Dhaval Sikligar, M.S., Computer Systems Architect, Tata Consultancy Services (TCS) – Industrial IoT & Low-Code Manufacturing Systems

Research Specialization & Interests:

Dhaval specializes in Industrial Internet of Things (IIoT) architectures, decision-support systems, and real-time manufacturing analytics. His key interests include driving operational efficiency through digital integration, leveraging low-code platforms (Tulip, Mendix) for rapid deployment, and embedding cybersecurity and data governance into industrial systems.

Professional Experience:

Over the past few years at TCS, Dhaval has architected and delivered critical digital solutions in the pharmaceutical, medical device, and broader manufacturing sectors. His flagship project involved designing and deploying a low-code IIoT decision-support system for a leading pharmaceutical manufacturer, integrating real-time data from Tulip Edge devices and MES platforms to enhance operational visibility and predictive risk management. As an IT Business Analyst, he also played a key role in streamlining operations—boosting stakeholder engagement by 30%, reducing system errors by 30%, and improving productivity by 15%. His contributions earned him multiple accolades, including TCS's Innovista Top Seeds and Applause Award in 2024.

In addition to his digital expertise, Dhaval brings robust field experience from his earlier roles in the core manufacturing sector. At Sealmatic India and Metcraft Steel Pvt. Ltd., he managed sales and operations for

industrial mechanical seal systems and stainless-steel infrastructure, respectively. There, he exceeded sales targets, led more than 20 successful tenders, implemented inventory management systems that reduced waste by 20%, and improved client retention by over 25%. This strong foundation in real-world production environments gives him a unique ability to translate operational challenges into scalable digital solutions that are both practical and transformative.

Contact: sikligardhaval@gmail.com

Table of Content

INTRODUCTION

Occupational safety, worker health protection, and environmental preservation—collectively referred to as Environment, Health & Safety (EHS)—remain core priorities for industry in the twenty-first century. Each year millions of occupational accidents and diseases still occur worldwide; according to the International Labour Organization, roughly 2.9 million people die annually from work-related causes [1]. Despite substantial progress in lowering injury rates—for example, in the United States the total recordable incident rate fell from 8.9 to 2.7 cases per 100 workers between 1992 and 2020—the number of fatal accidents declined far less sharply, by only 17 percent over the same period [2]. These figures highlight the limitations of conventional occupational-safety practices that rely mainly on reactive measures and minimal regulatory compliance. Contemporary conditions therefore demand a digital transformation of EHS systems that enables proactive risk management and assured compliance through seamless integration of data and technology.

Digitalisation has already reshaped many industrial sectors, demonstrating its potential to enhance both efficiency and safety. The COVID-19 pandemic accelerated this trend, underscoring the value of remote monitoring and automation for maintaining operational continuity [3]. Within EHS, new avenues have emerged: sensors and wearable devices for monitoring workplace conditions and worker status, industrial IoT platforms for tracking process parameters, and big-data analytics and machine-learning models for predicting emergencies [4]. Yet adoption often proves fragmented, lacking a unified approach that merges occupational safety, process safety, and environmental protection. Traditional EHS management systems—typically built around standards such as OSHA, EPA, GOST, and ISO—have evolved in parallel, giving rise to duplicated workflows and data silos [5]. Consequently, a pressing scientific and practical task is to develop an integrated digital EHS platform that aligns regulatory compliance with risk management within a single, cohesive framework.

The purpose of this monograph is to substantiate and develop an original methodology—"EHS-Digital Loop"—a digital loop for managing EHS compliance and risks at industrial enterprises, and to theoretically validate its effectiveness using secondary data. The central hypothesis posits that integrating disparate EHS systems into a unified end-to-end digital loop—

“Plan–Sense–Predict–Act”—will markedly improve safety and environmental performance by proactively identifying risks and ensuring timely fulfilment of requirements, while concurrently reducing costs through the elimination of redundant processes and the prevention of incidents.

To achieve this aim, the study addresses the following tasks: (1) to analyse the theoretical foundations and evolution of the EHS regulatory framework, identifying prerequisites for integration; (2) to design the architecture of the EHS-Digital Loop system, comprising modules for regulatory-requirements planning, sensor-data collection, predictive analytics, and automated response; (3) to define methodologies for generating leading safety indicators and establishing their relationship to lagging metrics; (4) to propose a unified meta-model for integrating the requirements of ISO standards (45001, 14001, etc.) and to ensure cybersecurity and data governance within such a system; (5) to conduct a theoretical piloting of the methodology across several scenarios (food, chemical, and construction sectors) using published data on the effectiveness of digital initiatives; and (6) to quantitatively assess the anticipated impact—in terms of ROI, carbon-footprint reduction, and resilience to stress factors—through analytical models and simulations.

The scientific novelty of this work lies in its comprehensive approach to the digital transformation of EHS. Whereas prior studies have focused on individual technologies—such as the application of IoT for occupational safety, big-data analysis of accident records, or integration of quality, environmental, and safety management systems—the present study proposes a single, original platform that unites regulatory and managerial aspects (compliance) with predictive and preventive mechanisms (risk management) within a continuous digital cycle. The practical significance is evident in the applicability of the results by industrial enterprises to enhance EHS system efficiency: scenario-based calculations indicate that implementing the EHS-Digital Loop can reduce injury and accident rates by 10–30% and improve environmental indicators. Moreover, the integrated system enables the avoidance of audit and documentation duplication—estimates suggest that a single audit of the integrated system can reduce labour efforts by up to 20% compared with separate audits [5, 6]. Thus, adoption of the proposed methodology delivers substantial economic benefits and fosters a culture of proactive safety.

CHAPTER 1: THEORETICAL AND REGULATORY FOUNDATIONS OF INTEGRATED EHS MANAGEMENT

1.1 Timeline of EHS Regulatory Development

EHS management systems for occupational safety, process safety and environmental protection have evolved in response to national and international legislation. To understand the prerequisites for integrating these domains, key milestones in the regulatory framework are summarised in Table 1.

Table 1. Key Milestones in the Development of the EHS Regulatory Framework

Year	Country/ Region	Regulatory Act or Standard	Significance and Impact
1970	United States	Occupational Safety and Health Act (OSH Act)	Established OSHA as the federal agency for workplace safety. Defined employer responsibilities for safe working conditions. Did not require formal management systems, but introduced standards and inspections.
1970	United States	Clean Air Act Amendments	First comprehensive federal law to limit air pollutant emissions. Empowered the EPA to set air quality standards. Spurred adoption of pollution-control technologies.
1972	United States	Clean Water Act	Regulated discharges of pollutants into surface waters. Introduced the National Pollutant Discharge Elimination System (NPDES) for industrial effluents. Laid the foundation for water-risk management at facilities.
1980s	United States	Expansion of EPA and OSHA Standards	New OSHA regulations on toxic substances (asbestos, lead, etc.) and EPA hazardous-waste programmes (RCRA, Superfund). Strengthened requirements for process safety.
1989	International	ILO Code of Practice No. 164: Occupational Safety, Health and the Working Environment	Provided recommendations for systematic workplace safety management. Anticipated later management-system standards.

Year	Country/ Region	Regulatory Act or Standard	Significance and Impact
1996	International	ISO 14001:1996 (Environmental Management Systems)	First international standard for environmental management. Required identification of environmental aspects, legal compliance and continual improvement. Became the global basis for corporate environmental programmes.
1999	International	OHSAS 18001:1999 (Occupational Health and Safety Management)	Developed by a consortium of certification bodies in the absence of an international safety standard. Required policy, risk assessment and incident-management processes. Widely adopted until ISO 45001.
2006	European Union	Regulation (EC) No 1907/2006 (REACH)	Comprehensive reform of chemical-substance regulation in the EU. Obliged manufacturers to register substances, assess risks and implement control measures. Industry assumed primary responsibility for chemical-risk management, often via integration into quality and environmental management systems.
2011	International	ISO 31000:2009/2018 (Risk Management)	Introduced a common approach to organisational risk management. Defined risk as "the effect of uncertainty on objectives," addressed both negative and positive effects, and established a universal process for risk assessment and treatment. Underpinned the risk-based approach in management-system standards (Annex SL).
2015	International	ISO 14001:2015 and ISO 9001:2015 (HLS via Annex SL)	Aligned EMS and QMS with a common High-Level Structure (HLS), sharing identical clauses (context, leadership, risk planning, support, operations, performance evaluation, improvement). Simplified the integration of different management systems and mandated the consideration of risks and opportunities in planning.
2018	International	ISO 45001:2018 (Occupational Health and Safety Management Systems)	First global standard for OH&S, replacing OHSAS 18001. Harmonised structure with ISO 9001/14001. Required hazard identification, risk and opportunity assessment, worker participation and continual improvement. Certification under ISO 45001 has been shown

Year	Country/ Region	Regulatory Act or Standard	Significance and Impact
			to improve both safety performance and organisational productivity.

As shown in Table 1, by the early 2020s the foundations for consolidating EHS systems were firmly in place. In the United States, the decades following the OSH Act of 1970 focused chiefly on establishing minimum standards and inspection-based enforcement. This approach drove a significant reduction in overall injury rates—for example, workplace fatalities declined from approximately 14 000 in 1970 to 5 250 in 2018. However, the OSH Act did not initially mandate enterprise-level safety management systems; emphasis remained on compliance with discrete regulations (such as hazardous-substance standards) and on responding to violations through penalties. Consequently, by the 2020s further improvements had plateaued: traditional measures no longer delivered the same returns, and experts have called for legislative updates and the adoption of modern EHS management approaches. Proposals include requiring companies to implement formal health-and-safety management programmes, extending OSH coverage to all worker categories, strengthening penalties for non-compliance, and updating regulations to reflect twenty-first-century realities [7].

In the realm of environmental regulation, the United States laws enacted between 1970 and 1972—the Clean Air Act and the Clean Water Act—laid the groundwork for systematic control of emissions and discharges. Over the ensuing fifty years, air and water quality have generally improved, although new challenges persist (for example, long-lasting PFAS chemicals and climate change). In its 2022 report marking the fiftieth anniversary of the Clean Water Act, the U.S. Government Accountability Office noted that the proportion of assessed water bodies has increased and that many rivers once heavily polluted have been restored. However, unmonitored sources (non-point discharges) remain problematic, and monitoring coverage is incomplete—only about fifty percent of U.S. waters have been evaluated [8]. This example illustrates that, even in the presence of robust legislation, effective risk management requires up-to-date data and an integrated approach.

In the European Union, the introduction of the REACH Regulation in 2006 represented a turning point: for the first time, the "no data, no market" principle obliged companies to conduct their own chemical-risk assessments and manage those risks [9]. REACH effectively unified occupational-health and environmental protection requirements for chemical substances by mandating consideration of their entire life cycle. Compliance with REACH spurred firms to implement chemical-safety management systems, often by extending existing ISO 14001 frameworks. Concurrently, since the 2010s the EU has promoted the concept of Integrated Management Systems (IMS), merging quality, environmental, and safety management. Guidance documents—such as the 2015 Annex SL for ISO standards—have simplified the creation of a single system. Recognised benefits of the integrated approach include elimination of duplication (unified procedures replacing three separate sets), reduced audit costs (for example, external-audit time declines by approximately twenty percent under a comprehensive IMS review [6]), and enhanced management efficiency and transparency [5].

By the mid-2020s, regulatory requirements—laws and standards alike—presume the existence of a systematic, risk-based EHS management system within enterprises. In practice, however, many companies continue to maintain separate, siloed functions: occupational safety in one department, environmental management in another, and process safety in yet another. At the same time, research demonstrates that organisations certified under integrated standards achieve superior outcomes. A recent analysis of 157 publicly traded companies found that ISO 45001 adoption is statistically linked to increased productivity and profitability compared with non-certified firms [10]. Other studies have reported lower injury rates at companies implementing OHSAS/ISO 45001 than at similar enterprises without certification [11]. These findings confirm the value of consolidating EHS systems into a single continuous-improvement loop, which is the focus of the present work.

1.2. Risk-Oriented Approach (ISO 31000)

The concept of risk-oriented thinking has become a fundamental element of EHS management, particularly following the publication of the international standard ISO 31000 "Risk Management – Guidelines" (first edition 2009, updated in 2018). According to ISO 31000, risk is defined as the "effect of uncertainty on objectives" [12]. This definition is intentionally neutral:

risk may produce both negative and positive outcomes. A positive outcome of uncertainty is identified in the standard as an opportunity, which should be considered alongside threats. Thus, risk management is a systematic process aimed at maximising positive deviations and minimising negative deviations from expected results. In the context of occupational safety and environmental protection, the primary focus naturally lies on negative risks (injuries, accidents, environmental damage), but opportunities (for example, adoption of a new safety-enhancing technology that yields benefits) are also taken into account.

A key element of the risk-oriented approach is the notion of risk acceptability. Since not all risks can be completely eliminated, it is necessary to establish what level of risk is deemed tolerable or acceptable for an organisation, taking into account regulatory requirements and its own criteria. In high-hazard industries, the ALARP principle (As Low As Reasonably Practicable) is applied, calling for risk reduction to the point where further mitigation would be technically or economically unjustified. Management-system standards (ISO 45001, ISO 14001) typically require procedures for risk and opportunity assessment, including the definition of acceptability criteria. For example, a company may determine that any risk carrying potential losses exceeding USD 1 million with a probability greater than 1 in 1 000 per year is unacceptable and demands mitigation, while risks below that threshold are considered acceptable under existing controls.

For risk visualisation and analysis, the bow-tie model is widely used (Figure 1)—a butterfly-shaped diagram at whose centre lies the undesired event (e.g. an accident), with causal factors and preventive barriers on the left and consequences along with mitigating barriers on the right.

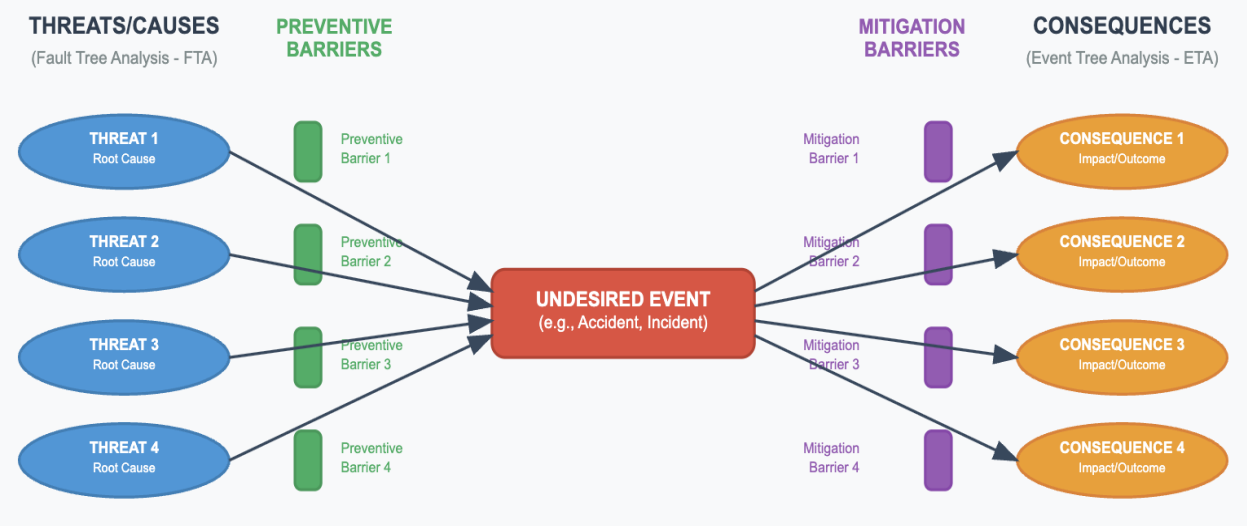

Figure 1. Bow-tie model

This model integrates causal analysis (fault-tree analysis, FTA) and consequence analysis (event-tree analysis, ETA) into a single representation [13]. Its main advantage is clarity: it shows how risk is controlled both at the prevention stage and during response. In process-safety management, the bow-tie method is extensively employed for complex hazards such as toxic releases, explosions, or falls from height (notably in the petrochemical and mining industries). In the contextual model of the EHS-Digital Loop, the bow-tie principle is likewise incorporated: during the planning phase, risk sources are identified and barriers (measures) established, while in the Act phase, actions to mitigate consequences—should the event occur—are defined.

Historically, risk assessment in EHS has often been conducted using matrix methods, where each risk is assigned a rank based on categories of likelihood and severity of consequences. For example, a 5×5 matrix (Figure 2) maps probabilities from “rare” to “frequent” against consequences from “negligible” to “catastrophic,” with colour zones (green, yellow, red) at each intersection indicating the risk level.

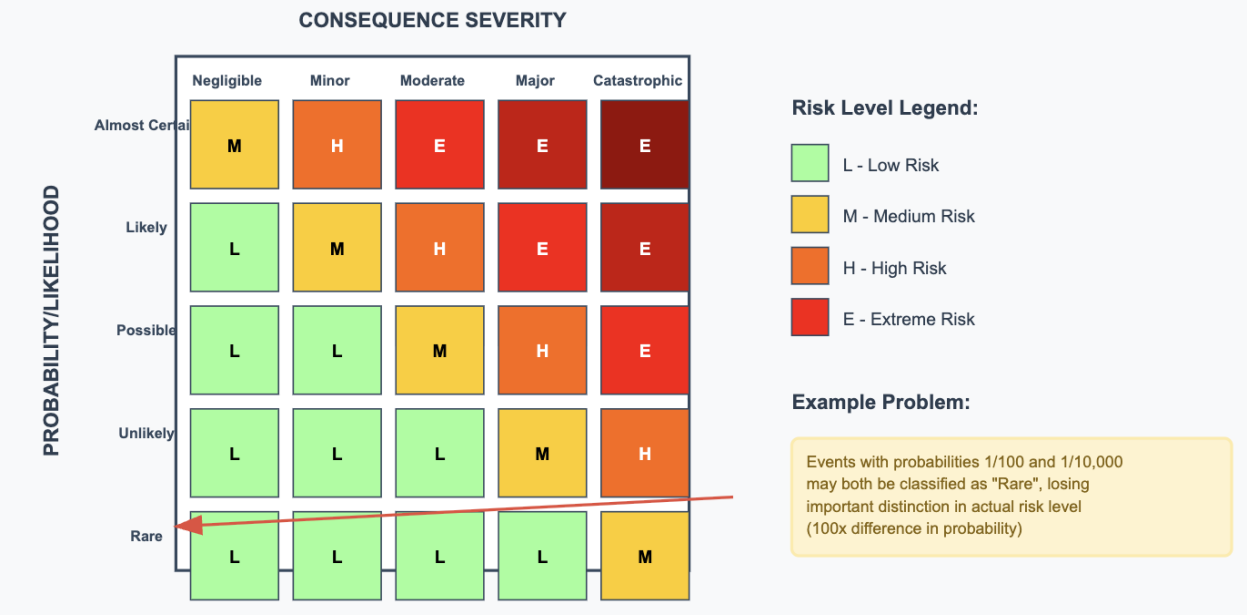

Figure 2. Traditional EHS Risk Assessment Matrix (5×5)

In this case, the matrix serves as a simple and intuitive tool, but it has significant limitations. First, matrix assessments are subjective: “likelihood” is often estimated qualitatively and may vary between experts. Second, the resolution of the matrix is low—aggregated categories conceal important nuances (for instance, events with probabilities of 1/100 and 1/10 000 may both be classified as “rare”). Third, matrices do not account well for interdependencies among risk factors and are ill-suited for dynamically updating risk levels as new data arrive [14].

With advances in technology, predictive risk scoring methods based on data analytics have emerged to overcome these matrix constraints. This approach involves gathering large datasets on incident precursors and applying machine-learning algorithms to estimate both event probabilities and factor importance. For example, in the construction sector, researchers compiled data on hundreds of projects—project characteristics and internal inspection outcomes—and employed classification models (Random Forest), achieving approximately 78 percent accuracy in predicting accidents. Another study uncovered cyclical patterns: immediately after an injury, a firm intensifies preventive measures (briefings, observations), which later diminish—only for accident probability to rise again over time [4]. Such insights cannot be derived from a static risk matrix but follow from time-resolved data analysis.

It should be noted that the adoption of Predictive Analytics in EHS remains nascent yet is rapidly evolving. Several barriers persist: sufficient volumes of high-quality data are not always

available; EHS professionals require data-analysis expertise; and model outputs must be communicated in clear, actionable terms. The last aspect—model interpretability—is critically important. Modern techniques such as SHAP (Shapley Additive Explanations) allow the contribution of each feature to a model's prediction to be quantified [15, 16]. Thus, even a "black-box" method like gradient boosting can be made transparent—for instance, if a model forecasts a high accident probability for a particular division, SHAP can reveal that the largest contributors were factors such as monthly overtime exceeding 20 hours and a high count of recorded minor incidents (near misses). In the proposed approach (see Chapter 2), explainable models are embedded so that predictive results can be communicated directly to managers and auditors for timely action.

Accordingly, the modern risk-oriented approach in EHS encompasses systematic hazard identification; risk assessment that considers both probability and severity; establishment of acceptability criteria; planning of risk responses (elimination, reduction, transfer, acceptance); and continuous monitoring and revision of risks based on empirical data. Integrating this approach across all organisational processes—from strategic planning to operational tasks—is mandated by new ISO standards and recognised as best practice. Chapters 2–4 will illustrate how the EHS-Digital Loop platform implements the risk-oriented approach in practice, combining classical principles (e.g., bow-tie, ALARP) with big-data and AI capabilities to enhance the accuracy and timeliness of risk assessment.

1.3. Current Challenges

Before outlining the proposed methodology, it is essential to identify the bottlenecks and issues common to many existing EHS management systems. These challenges define the targets for improvement that digital integration must address.

1.3.1. Data Silos and Insufficient End-to-End Analytics

In many organisations, occupational-safety, process-safety and environmental-protection data reside in separate systems. For example, the OHS department maintains its incident and injury log in one database, the environmental team records emissions and waste reports in another, and the production-control unit uses a third. As a result, a consolidated, enterprise-level view of risk is often absent—the data remain unlinked and are never analysed holistically.

A NIOSH study observes that although companies amass ever larger volumes of data, a substantial share goes unanalyzed and unused for injury prevention [4]. Traditional EHS information systems (such as local logs or Excel registers) lack the capacity to merge data from disparate sources. Consequently, organisations overlook valuable correlations—for instance, the relationship between production disruptions and personnel incidents, or between environmental breaches and safety violations.

The absence of a unified data repository also hinders the deployment of advanced analytics tools. Addressing this requires integrating all EHS data onto a common platform and fostering a data-driven decision-making culture in safety management. Recognising the need, the U.S. National Safety Council (NSC) launched its Work-to-Zero programme, which leverages data analytics to identify technological solutions for reducing injury rates [2]. The EHS-Digital Loop solution envisages creating a single EHS data lake, into which all relevant information—from audit results to sensor readings—flows for comprehensive, end-to-end analysis (see Chapter 2).

1.3.2. Focus on Lagging Indicators and Reactivity

Traditional EHS KPI systems are built around lagging indicators—metrics such as the lost-time injury rate (LTIR), total recordable incident rate (TRIR), and number of days without accidents. These measures reflect events that have already occurred, essentially quantifying past unsafety.

Organizations often set targets like “achieve an LTIR below X” or “zero injuries,” monitor these metrics, and report them in annual reviews [11]. However, this approach carries an undesirable side effect: attention shifts to the numbers rather than to root causes. Hitting a zero-injury target may foster complacency, even if numerous no-lost-time incidents or hazardous conditions occurred but, by chance, did not result in serious outcomes.

Emphasis on lagging KPIs can also encourage underreporting of minor incidents to “keep the statistics clean.” In both academic literature and leading practice, there is growing focus on leading indicators—measures of preventive activity such as the number of hazards identified and rectified, frequency of safety-observation audits, or workforce fatigue levels as tracked by wearables [12].

Leading indicators make it possible to assess system health before incidents occur and to trigger proactive measures. The challenge lies in selecting and measuring these indicators: they are often industry-specific and lack universally accepted benchmarks. Nevertheless, the shift from a purely reactive mindset ("investigating incidents") to a proactive one ("managing precursors") is recognised as essential worldwide [17, 18]. Many organizations begin to implement leading metrics, but often do so manually—via surveys or observations—without systematic analytics to link them to outcomes.

The proposed approach (see Section 2.2) aims to close this gap. It leverages IoT data and log records to automate the creation of leading indicators—such as trends in worker fatigue metrics or frequency of sensor threshold breaches—and correlates them with traditional KPIs, thereby enriching the risk-management framework [4]. Refocusing on leading indicators should reduce the lag between risk escalation and organizational response.

1.3.3. Duplication of Audits and Disparate Compliance Processes

In organisations where individual EHS functions operate independently, the same tasks are frequently repeated. For example, the environmental-protection department conducts its annual audit for compliance with environmental legislation, while the occupational-safety service carries out a separate safety audit. Although both may examine overlapping areas (such as chemical storage—critical for both environmental protection and worker safety), they do so in isolation. Documentation is likewise duplicated: parallel registers of regulatory requirements are maintained—one for environmental, another for occupational safety—even though many requirements intersect (for instance, control of volatile-organic emissions is both an environmental mandate and a critical ventilation requirement for worker safety).

Standards and regulations often do not align, necessitating multi-step verifications. A facility might undergo external certification audits—one for ISO 9001, another for ISO 14001, and a third for ISO 45001—each scheduled in different months. This disperses personnel resources, as similar documentation must be prepared three times for different auditors.

International experience demonstrates that integrated management systems optimise resources: according to the International Accreditation Forum (IAF), a combined IMS audit typically requires 10–20 percent less time than the sum of separate audits [6]. Moreover, a single

system eliminates contradictions: instead of three distinct policies (quality, environment, safety), one unified IMS policy is established [5], with all processes aligned under its umbrella.

It must be emphasised that duplication is not merely an extra cost but also a source of inefficiency: when responsibility is diffused across departments, some issues can "fall between the cracks." For instance, a process-safety audit may overlook an environmental-risk aspect (if deemed outside its scope), and vice versa. An integrated audit or process-based review allows for a comprehensive assessment of risk. Consequently, our methodology (see Chapters 2.3 and 2.5) prioritises the creation of a unified registry of obligations and risks and the holistic monitoring of requirement fulfilment, so that a single tool covers all EHS dimensions. This approach will eliminate the need for parallel spreadsheets and separate reviews—one system will automatically track compliance across every domain and generate consolidated reporting.

1.3.4. Vulnerability of Operational Technologies (OT) in EHS Digitalization

The deployment of sensors, Industrial Internet of Things (IIoT) devices, and the connection of production equipment to networks carries a downside—cyber risks. Historically, automation systems (OT/ICS) at industrial sites were air-gapped (control networks were not linked to corporate or external Internet). In the Industry 4.0 era, however, an increasing number of devices support network connectivity and open protocols (for example, many controllers now use OPC UA or MQTT to stream data to external analytics platforms—see Chapter 2.1). While this enables real-time EHS monitoring, it also introduces new attack surfaces.

There have been documented incidents in which malware targeted industrial-safety systems—for example, the Triton/Trisis attack in 2017, aimed at safety-instrumented-system controllers at a chemical facility. Adversaries might disable sensors that inform risk assessments or disrupt alarm and notification systems. Industry reports indicate that cybersecurity incidents in manufacturing and process industries have increased in recent years [19].

Within the EHS context, these developments mean that any digital platform must be resilient to cyber threats; otherwise, it risks creating a new category of hazard. He et al. (2023) identify a range of technical challenges for IoT in high-hazard industries, including sensor energy efficiency and reliability, network scalability, standardization and interoperability, as well as security concerns (data privacy, channel integrity, protection against unauthorized access). An

expert survey highlighted 28 key IoT challenges for EHS, with the greatest weight given to limited community-level support (safety culture), standardization gaps, and issues of data reliability and privacy [19].

In this monograph (see Chapter 2.4), cyber vulnerabilities are addressed through Zero Trust Architecture principles—where no device is trusted by default and access is granted under a least-privilege model—and by adhering to the IEC 62443 series of industrial-cybersecurity standards. Compliance with IEC 62443 entails network segmentation, strict device authentication and authorization, data encryption, traffic-anomaly monitoring, and more. Special attention is paid to privacy protection: wearable devices collect personal worker data (e.g. heart rate, location), and any leakage or misuse is unacceptable. Accordingly, the platform stores personal data in pseudonymised form and ensures GDPR compliance for all records (for example, obtaining employee consent and using data only in aggregate analyses) [20].

In summary, the current shortcomings of many EHS systems can be reframed as opportunities for enhancement via digital technologies and integration. Transforming fragmented data into connected insights through a unified platform, shifting the focus from reaction to prevention via analytics, optimising EHS administration through integrated processes, and ensuring cyber resilience—all these elements are embodied in the EHS-Digital Loop concept presented in the next chapter.

1.4. Baseline Set of Industry KPIs

For the purposes of modelling and evaluating the effectiveness of an integrated digital system, it is necessary to establish a baseline level of key EHS indicators against which improvements will be measured. Since the theoretical validation of this study (Chapters 3–4) does not rely on data from a specific enterprise but aggregates information from literature and reporting, it is reasonable to adopt industry-average metrics for injury rates and environmental performance. These figures will serve as reference points: representing the “before digital integration” state, against which the “after” scenario can be modelled.

The most widely used metrics for workplace injuries are:

- LTIR (Lost Time Injury Rate) – the rate of injuries resulting in lost work time, typically calculated as the number of such injuries per 200 000 work-hours (equivalent to 100

full-time employees over one year). In some sources it is also referred to as LTIFR (Lost Time Injury Frequency Rate).

- TRIR (Total Recordable Incident Rate) – the total rate of recordable injuries and incidents (including all cases requiring medical treatment, with or without lost time), likewise per 200 000 work-hours.

Different countries use analogous measures—for example, in the EU the frequency rate may be expressed per 1 000 000 hours, or elsewhere per 1 000 employees. For comparability, this work adopts the per-200 000-hours metric common in the United States, since many international firms publish their data in this format.

According to the U.S. Bureau of Labor Statistics, the aggregate injury rate in the private sector in recent years has been approximately 2.6 cases per 100 full-time equivalents (i.e. TRIR ≈ 2.6) [21]. For example, in 2018 the TRIR stood at 2.8, and in 2023 at 2.6 (Figure 3) [21]. These figures represent an average across all industries. Accordingly, TRIR ≈ 2.6 (per 200 000 hours) can be adopted as the baseline level for large-scale industry. Highly hazardous sectors exhibit higher rates (such as construction and mining), while office environments typically report lower values.

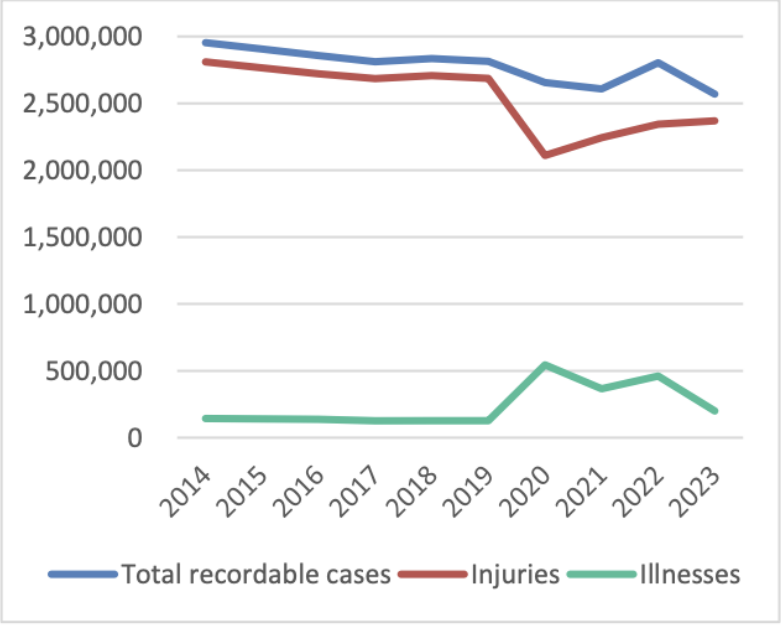

Figure 3. Total recordable injury and illness case counts, private industry, 2014–23 [21]

With respect to LTIR, this metric is naturally lower than TRIR—since not all recordable incidents result in lost work time—and its average values vary by industry. For example, the worldsteel association reported an average LTIFR of approximately 0.76 per 1 000 000 hours in 2024, equivalent to about 0.15 per 200 000 hours [22]. In contrast, manufacturing and food-processing sectors typically report values closer to 1.0. ESG disclosures from many multinational food, beverage, and consumer-goods companies often cite LTIRs in the range of 0.5–1.5 per 200 000 hours. For modelling purposes, a baseline LTIR of 1.0 is adopted—recognised as a conservative average, with 0.5–1.0 representing a range of good practice. This corresponds to one lost-time injury per 200 000 work-hours (roughly one per 100 employees per year), a level characteristic of relatively safe large enterprises.

Beyond occupational-safety metrics, carbon intensity of production serves as a critical integrative KPI for environmental performance. Expressing emissions per unit of output or energy, it links operational-efficiency gains to climate-impact reductions. To standardise this indicator, CO_2 emissions per kilowatt-hour of electricity generation are employed as a proxy for equivalent emissions savings. The International Energy Agency's 2021 data place the global, weighted-average emission factor at about 0.475 kg CO_2/kWh [23], although regional values vary—lower in areas with low-carbon generation (e.g. Europe) and higher where coal predominates. The task reference cites a coefficient of 0.233 kg CO_2/kWh, reflecting updated global or regional averages circa 2024. Indeed, IEA Emission Factors 2024 data show declining carbon intensity in many countries, with the OECD average near 0.25 kg/kWh [24]. For simplicity, 0.233 kg CO_2/kWh is taken as the baseline conversion factor for energy-savings-to-emissions-reduction modelling.

Finally, typical direct-cost estimates for incidents are required for ROI calculations. Published figures vary: the U.S. National Safety Council assesses the average economic loss from a serious lost-time injury at around USD 54 000—including medical expenses, compensation, and productivity losses [2]—while large corporations report even higher costs (up to USD 100 000–150 000 when indirect impacts are included). The subsequent analysis will employ multiple scenarios, with a base-case direct cost of USD 50 000 and sensitivity bounds of ±20 % around that value.

Thus, the baseline set of KPIs for use in the calculations is defined as follows:

- TRIR (Total Recordable Incident Rate): 2.6 cases per 200 000 hours – the averaged industry-level rate .
- LTIR (Lost-Time Injury Rate): 1.0 case per 200 000 hours – an average value (in practice, industry leaders report lower figures, while some sectors exceed this; selected for scenario modelling) .
- Energy Carbon Intensity: 0.233 kg CO_2/kWh – the average emission factor reflecting the contemporary energy mix .
- Average Cost of a Single Serious Incident: USD 50 000 (direct and indirect losses; approximate NSC estimate) .
- Annual Energy Throughput per Facility: for example, 100 GWh/year (representative of a large-scale operation; used illustratively in the calculations and scalable as required).

This baseline will be employed in Chapter 3 for "before/after" digitalisation comparisons across case studies, and in Chapter 4 for constructing the economic-environmental models. In particular, the scenarios include:

- Food Sector: modelling an LTIR reduction from 0.9 to 0.6 (aligned with our baseline of ~1.0).
- Chemical Sector: using a PSER of 0.3 (incident rate per 200 000 h for process operations, to be mapped onto LTIR).
- Construction Sector: starting from a TRIR of ~3.5 (typically above average) and simulating potential reductions, etc.

Normalisation of all metrics to a common denominator (per 1 000 000 hours) will be conducted in Section 3.4.

It is important to emphasise that, while these averaged indicators do not capture the specifics of any single enterprise, their application is justified for the theoretical illustration of trends. As the objective is to demonstrate relative improvements and model viability, the precise absolute values are less critical; nevertheless, their grounding in empirical statistics enhances the credibility of the findings.

CHAPTER 2: THE EHS-DIGITAL LOOP METHODOLOGY

2.1 System Architecture ("Plan → Sense → Predict → Act")

The proposed methodology is structured as a closed digital management loop comprising four phases: Plan, Sense, Predict, and Act. These stages form a continuous cycle that drives ongoing enhancement of the EHS system—analogous to Deming's PDCA cycle, but leveraging digital monitoring and analytics capabilities. Figure 4 schematically depicts the interplay among the phases of the EHS-Digital Loop cycle.

Figure 4. Cyclic architecture of the EHS-Digital Loop: the four phases of the integrated EHS management process.

Each phase is supported by dedicated functional modules within the system (see Sections 2.1.1–2.1.4).

2.1.1. Phase 1: Plan (Planning)

At this stage, the system automates both the maintenance of a register of regulatory requirements and internal standards, and the processes of hazard identification and risk assessment.

In the traditional approach, specialists periodically monitor legislative changes, update requirement lists, and conduct risk assessments (for example, reviewing the enterprise risk register once a year). In the EHS-Digital Loop, this work is supported by a dedicated module: a digital register of requirements linked to external sources (legislative and standards databases). This register can update itself automatically—when a new regulation or amendment is issued, the system receives a notification via API from regulatory-information providers [9] and adds the new requirement to its listing. Each requirement is assigned attributes such as domain (occupational safety, environmental protection, etc.), applicability to the specific enterprise, compliance deadlines, and responsible parties.

Additionally, during the Plan phase a hazard register and associated risk assessments are compiled: based on the enterprise's operations, all significant hazards (process-related, chemical, ergonomic, environmental, and so on) are listed, and a risk assessment is carried out for each—initially expert-driven, indicating probability and severity or, where available, drawing on existing data. This hazard register is then integrated with the requirements register; for example, if an "ammonia leak" hazard is identified, the system links it to the relevant regulations (safety rules for refrigeration systems, maximum permissible concentrations, and the like).

Thus, the Plan module creates the knowledge base for the entire system—defining what must be done and which risks demand attention. Priorities are set either by experts (for example, assigning risk levels via a matrix) or by leveraging data, as described in Section 2.2. The outputs of the Plan phase are: (a) a digital regulatory register (in the form of a database or JSON schema, see Section 2.3) and (b) a ranked list of risks with specified control measures. These Plan-phase data serve as the inputs for the Sense and Predict phases, guiding the system on what to monitor and what to anticipate.

2.1.2. Phase 2: Sense (Data Collection / "Sensing")

This phase is devoted to continuous monitoring of working conditions, behaviour, and the surrounding environment using sensors and other systems. The Industrial IoT paradigm calls for equipping facilities with as many sensors as possible—both fixed (inline flow meters, gas analyzers, cameras) and wearable devices for personnel (smart helmets, trackers, fatigue sensors). In the context of EHS-Digital Loop, the Sense module unifies data streams from multiple sources:

- Worker-safety IoT sensors. Wearable devices monitor body posture (to detect falls or improper lifting form), heart rate and fatigue levels, and even include "panic buttons" on bracelets for emergency calls. Today, devices exist that detect micro-sleeps in drivers or elevated carbon-monoxide levels in welders [20, 26]. These data (timestamp, worker ID, measured value) are transmitted in real time over wireless networks (Wi-Fi, Bluetooth, LPWAN, depending on coverage).
- Fixed environmental and equipment sensors. These track parameters such as microclimate conditions (temperature, humidity), concentrations of hazardous gases (H_2S, CO, VOCs, etc.), noise levels, and vibration. Equipment-health sensors (vibration, pressure, rotation speed) can signal deviations indicative of impending failures. Many modern industrial instruments support the OPC UA protocol (Open Platform Communications Unified Architecture), an open standard that enables "machine-to-machine" integration across vendors [4]. Another common protocol for telemetry is MQTT (Message Queuing Telemetry Transport), a lightweight TCP/IP-based protocol optimised for streaming sensor data to a central platform. In our architecture, all sensors register with an MQTT broker, and the EHS-Digital Loop core subscribes to relevant topics (e.g. "safety/worker1/heartbeat" or "env/zoneA/H_2S") to receive real-time updates.
- Maintenance-management systems (CMMS) and near-miss reports. Beyond physical sensors, existing corporate systems provide invaluable EHS data. A CMMS records equipment failures, work orders, and inspection results—indirect indicators of safety status, since frequent breakdowns may signal heightened accident risk. A time-tracking system can reveal overtime patterns, a known risk factor for injuries [2]. Finally, a near-miss registry—logging events that almost resulted in harm—is the gold mine of proactive safety. Within EHS-Digital

Loop, these systems are integrated via APIs or data uploads, creating a unified event stream. For example, each logged near miss (e.g. “operator slipped but recovered balance”) flows into the Sense module as an input for predictive models.

Thus, the Sense phase delivers a comprehensive situational picture: the system “feels” everything of consequence on site—from equipment parameters to each worker’s status. To handle this data volume, preprocessing rules and algorithms are configured. For instance, instead of storing all 86 400 per-second temperature readings, the system may compute rolling averages or record only out-of-range events. Edge modules connected to the MQTT broker can perform local aggregation. Overall, the Sense module constructs an up-to-date digital twin of the EHS environment: a timestamped, geo-referenced stream of parameters and events.

2.1.3. Phase 3: Predict (Forecasting)

This phase constitutes the core of the intelligent analytics within the EHS-Digital Loop. During Predict, the system processes incoming data from the Sense module using machine-learning and statistical techniques to uncover latent patterns and forecast likely incidents. Key tools include:

1. Supervised learning on historical data. When a repository of past events (injuries, accidents) exists, models can be trained to predict incident probability based on preceding factors. For example, using multi-year incident records, the system might train a gradient-boosting or neural-network model with inputs such as department, season, average worker tenure, number of violations detected in the prior month, overtime hours logged, and sensor readings (temperature, equipment load, etc.), and with the output being whether an incident occurred (yes/no) in the following month.

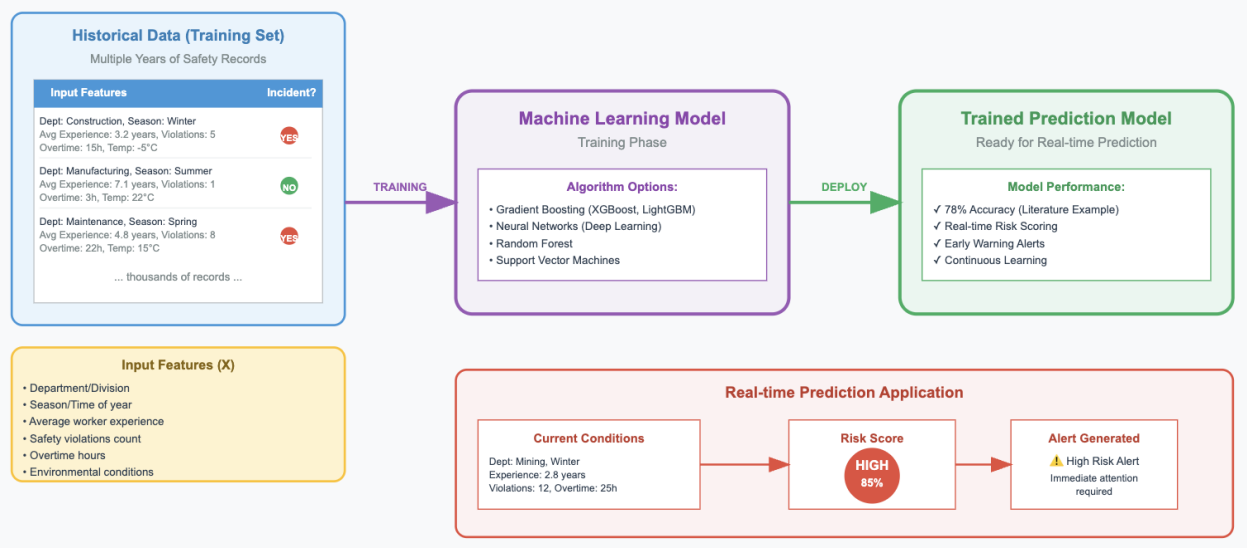

Figure 5. Supervised Learning for EHS Incident Prediction

Once trained, this model operates in real time: a high risk score for a given area generates an immediate alert. Studies have demonstrated the feasibility of such approaches—for instance, the aforementioned construction-site model achieved 78 percent accuracy in distinguishing between accident-prone and safe periods [4]. In our platform, similar models are initially trained on external datasets and then refined using organisation-specific data. Particular emphasis is placed on forecasting based on precursor events—small leaks or minor injuries—as signals of potential major incidents. Approaches such as "Safety II" recommend analysing not only failures but also successful interventions (what went right), and these positive outcomes can also feed into the model as resilience indicators.

2. Analysis of worker behaviour and condition. Wearables enable the prevention of human-factor injuries—fatigue, inattention, and the like. Smart cameras and bracelets can detect signs of exhaustion (sluggish movements, prolonged immobility) or physiological stress (elevated heart rate and temperature indicating possible heat stress). Based on these indicators, the system can forecast a high likelihood of error or loss of consciousness in the near term and issue a warning. (Although the alert action itself occurs in Phase 4, the decision logic originates here in Predict.)

3. Impact assessment of changes and scenarios. The Predict module also functions as a "what-if" simulator. If a change in process or replacement of equipment is planned, the system

can evaluate its effect on the risk profile using embedded risk models. This operates like an EHS digital twin: for example, the model may simulate that a 20 percent increase in load on Line X raises the probability of Pump Y failure by Z percent, thereby increasing accident risk.

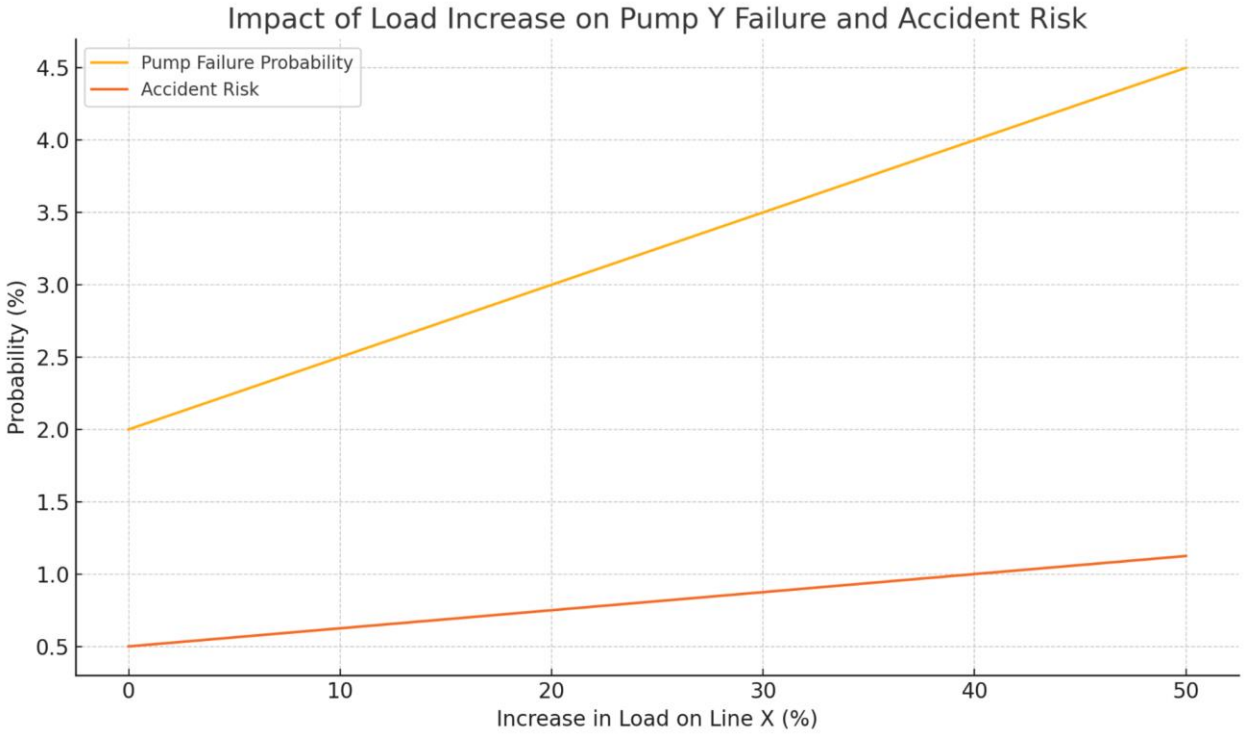

Figure 6. Impact of Load Increase on Pump Y Failure and Accident Risk

4. Explainable AI for trust and auditability. Forecast outputs must be transparent to safety engineers and management. Techniques such as SHAP (Shapley Additive Explanations) or similar methods quantify each input feature's contribution to a given prediction [16, 27]. For instance, if the system assigns a "High" risk level to Workshop 5 for the next day, SHAP can decompose the score—"20 percent due to three newly onboarded workers lacking full onboarding," "15 percent due to forecasted temperatures above 30 °C (heat wave)," "10 percent due to overdue maintenance on equipment," and so on. This breakdown helps the responsible manager understand the drivers of elevated risk and the control measures likely to mitigate it—serving as a bridge to the Act phase. Explainable-AI capabilities are especially critical for external audits: if the algorithm recommends actions based on a complex neural network, auditors may request justification. Our system will generate these explanations automatically, enhancing transparency and user trust [16].

Thus, Predict transforms raw Sense data into actionable forecasts and warnings. Whereas traditional systems relied on an engineer's intuition ("many fatigued workers mean expect an incident"), here AI scales and strengthens that role by detecting complex correlations (for example, the combination of factors A, B, and C persisting for three days yields an 80 percent chance of Incident D). It is important to emphasise that the model does not replace human judgement but augments it: the final decision to act remains with a qualified professional. However, Predict equips stakeholders with early warnings that, in conventional practice, would only emerge post-incident.

2.1.4. Phase 4: Act (Action)

The final stage of the cycle is the implementation of responsive measures for identified risks and non-conformities. Here, the digital system moves from analytics to active intervention in processes by generating tasks, notifications, and automated commands. Within the EHS-Digital Loop, the Act module comprises several tiers of action.

First, directives and alerts are issued to responsible personnel. If Predict has generated a critical forecast or detected a compliance breach, the system creates an appropriate notification. Examples include:

- "Zone A: equipment overheating likely tomorrow – schedule an unscheduled inspection."
- "Requirement: grounding check – due tomorrow and not yet completed!"

These alerts are addressed to the relevant roles—shop-floor manager, EHS engineer, department head—depending on the nature of the risk. Notification channels may include email, push notifications in a mobile app, SMS, and dashboard displays. The system can also alert workers on site—for instance, vibrating a wrist-worn device to warn an individual approaching a hazardous zone (based on real-time location data processed in Sense). Many modern smart helmets and vests have visual or audible alarms; our platform can command these devices to trigger a general evacuation alarm if, for example, a gas concentration threshold is exceeded.

The next tier consists of rule-driven responses based on standards: a distinctive feature of the EHS-Digital Loop is its rule engine, a set of mappings that link predicted events or detected

violations to prescribed actions. These rules derive from both internal procedures and external regulations. For example:

- If a heat-stroke risk is forecast for a worker, the procedure mandates moving them to a cool area and supplying water.
- If a chemical leak is detected, the emergency plan requires shutting down the pump and activating ventilation.
- If the system notes that ISO 45001 clause 6.1 ("Actions to address risks and opportunities") is overdue, it will generate a corrective-action task.

Rules can be organized in a correspondence table (Table 2). In many respects, this rule engine resembles an incident-management system (IMS), but here it operates automatically, driven by incoming data and forecasts.

Table 2. Situation – Action – Responsible – Deadline

Situation	Action	Responsible	Deadline
Forecasted heat stroke for a worker	Move to a cool area and provide drinking water	Safety Officer	Immediately
Detected chemical leak	Shut down pump and activate ventilation	EHS Technician	Immediately
Overdue ISO 45001 clause 6.1 ("Risk & Opportunity Management")	Generate and assign a corrective action	Quality Manager	Within 24 h
Fall risk above threshold	Inspect and secure safety harness	Site Supervisor	Immediately

The next level is the automatic initiation of tasks and work orders via the mobile platform. One of the outcomes of the system's action can be the creation of an electronic work order for personnel. Modern EHS platforms often include an event-management module: when a non-conformity is detected, the responsible manager assigns a task and the assignees receive it in a mobile app, then confirm completion [2].

In our implementation, this chain is partially automated: the system itself will generate the task, assign it to the appropriate person according to pre-configured matrices (for example, a

mechanic for equipment issues, an environmental engineer for ecology-related items, etc.). The assignee receives a notification on a smartphone or tablet (many workers today are issued corporate devices for briefings and communication). For example:

> **Task:** Conduct an unscheduled toolbox talk with Crew X before today's overtime shift (system detected elevated risk due to fatigue).
> **Assigned to:** Shift Supervisor
> **Due by:** 18:00

Upon completion, the assignee marks the task done—optionally attaching a photo or comment—and that information returns to the system (closing the loop by updating the Plan phase's risk register).

Finally, there is limited direct control of engineering systems. In certain cases—if pre-approved and programmed—the platform can issue commands to equipment without human intervention. For instance, if predictive-analytics sensors detect an unacceptable rise in compressor vibration (indicating imminent failure), the system could automatically issue a safe-shutdown command via the SCADA system. Such functionality must be deployed cautiously, as an automatic shutdown can itself trigger disruptions; in early stages, the system may instead issue "shutdown recommended" prompts that an operator must confirm. Nevertheless, EHS-Digital Loop is designed for integration with process-control systems via open APIs. The ANSI/ISA-95 standard (enterprise-to-control system integration) supports these integration scenarios, and our approach is fully compatible with it.

Through these capabilities, the Act phase ensures the cycle is closed: insights and analytics become real-world actions that enhance EHS performance. All completed actions are logged and automatically update the database: when a hazard is eliminated, it is removed from the next risk review; when a requirement is fulfilled, the regulatory register is updated (no further non-conformity remains).

Thus the cycle recommences: Plan ingests the updated inputs from Act, Sense continues monitoring under the new conditions, Predict recalibrates to the revised situation, and Act once again implements measures—ad infinitum. Conceptually, the EHS-Digital Loop establishes a system striving for maximal incident prevention—ideally zero accidents (the "Vision Zero"

initiative) [2]. While complete elimination of risk is unattainable in practice, this proactive loop greatly improves the likelihood of approaching that goal.

It is important to note that this architectural approach aligns with current management trends. It melds the ISO PDCA cycle (Plan-Do-Check-Act), the military OODA loop (Observe-Orient-Decide-Act), and the IIoT paradigm (data capture – analysis – action) in industry. We have deliberately adopted the "Plan-Sense-Predict-Act" terminology to highlight digital sensing and forecasting as new elements absent in classic PDCA. In the following sections of Chapter 2, each module's implementation details will be explored in depth. Overall, the EHS-Digital Loop architecture delivers a systematic, continuous EHS management process in which every requirement is tracked, every risk monitored, every significant signal analyzed, and every deviation triggers corrective action. This holistic loop forms the foundation for achieving a qualitatively higher level of safety and environmental performance within the organisation.

2.2. Analytics of Leading Indicators

As noted in Section 1.3, one of the principal shortcomings of traditional EHS systems is their reliance on lagging indicators—injuries, incidents, and losses. To shift toward proactive prevention, it is essential to establish a suite of leading indicators that signal rising risk before a serious event occurs. In the EHS-Digital Loop methodology, leading-indicator analytics occupy a central role, linking the Plan, Sense, and Predict phases: during planning, potential precursors are identified (what to track); in the Sense phase, those metrics are measured; and in Predict, their statistical relationship to outcome events is used to assess risk levels.

Drawing on both practice and literature, leading indicators in occupational safety and environmental protection can be grouped into several categories [4]:

- Activity metrics. These directly reflect preventive efforts: for example, the number of safety toolbox talks held before shifts, the count of hazardous conditions identified during site walks, or the percentage of planned safety actions completed. A decline in these metrics often precedes an uptick in injuries, as attention to safety wanes [4]. In the environmental domain, comparable measures include the number of completed process-control checks or waste-treatment equipment inspections.

- Behavioural indicators and safety culture. Metrics such as the proportion of workers submitting near-miss reports (staff engagement), or the results of peer-observation audits, fall into this category. Studies show that a strong reporting culture—where minor hazards are actively logged—is correlated with fewer serious incidents [4]. For instance, 50 near-miss reports in a month signal vigilant hazard recognition, whereas zero reports may indicate under-reporting or hidden risks.
- Equipment- and process-health indicators (accident precursors). Modern machinery is fitted with vibration, temperature, pressure, and electrical sensors. Tracking these signals can reveal early degradation: a 10 percent vibration increase over a month, a 5 °C temperature rise above baseline, or a surge in minor breakdowns. Process-safety leading indicators might include the number of safety-valve actuations, detected small leaks, or the share of equipment that has reached its service life without replacement. Worsening trends warn of imminent major failures.
- Worker health and fatigue indicators. Wearables and physiological sensors enable monitoring of average fatigue levels (for example, via sleep-stage tracking devices or fatigue surveys), on-task heart-rate trends (persistent elevation signals overload), and overtime hours logged. Accumulated fatigue and stress are proven contributors to errors and injuries [2]. Thus, metrics like "average overtime hours per worker in the past week" or "percentage of staff working over 12 consecutive days without rest" serve as critical leading indicators—high values portend an increased incident rate.
- Process-deviation indicators. Examples include the frequency of equipment entering alarm or emergency modes (even without a consequential shutdown), the number of bypassed alarms, and the percentage of product parameters falling outside specification (scrap rates, raw-material overuse). These indirect metrics flag process instability that can culminate in accidents.

In the EHS-Digital Loop, the Plan phase begins by selecting and configuring the leading indicators most sensitive to the risks of the specific enterprise. The platform provides a comprehensive builder: any connected system—from mobile audit apps to SCADA or wearable sensors—can serve as a source of numeric values, which are then calculated, stored, and visualised automatically.

As an example, consider the following indicators:

Leading Indicator 1 – “Number of safety-requirement violations detected per week.”. When the Weekly Safety Toolbox Talks line falls below the corporate minimum (see Fig. 7), the count of recorded violations almost invariably rises. That inflection point in the curve acts as a trigger: the rule engine generates a corrective task for the HSE manager before injury statistics begin to climb.

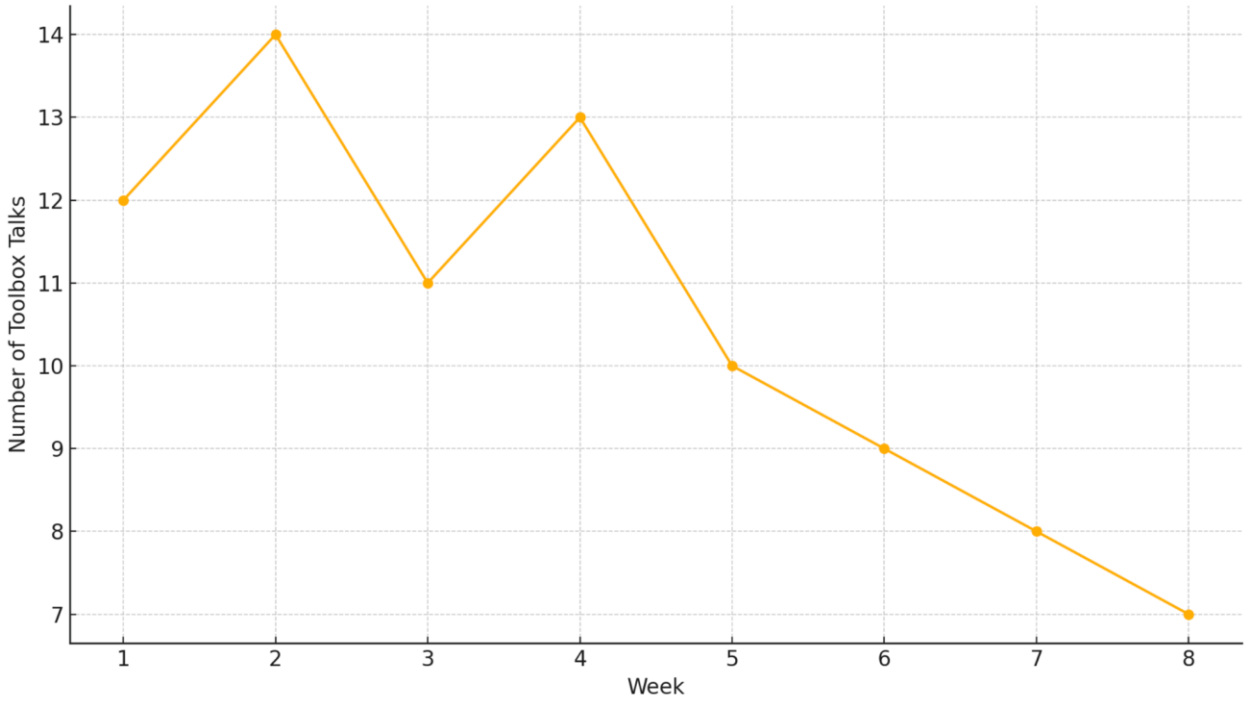

Figure 7. Activity Metric: Weekly Safety Toolbox Talks

Leading Indicator 2 – “Percentage of employees who have not completed the mandatory safety briefing on time.” A drop in near-miss reports (Fig. 8) coupled with a simultaneous increase in overtime hours (Fig. 9) signals workforce fatigue and a weakening safety culture. Once the lateness threshold is exceeded, the rule engine schedules a refresher briefing and notifies the shift supervisor.

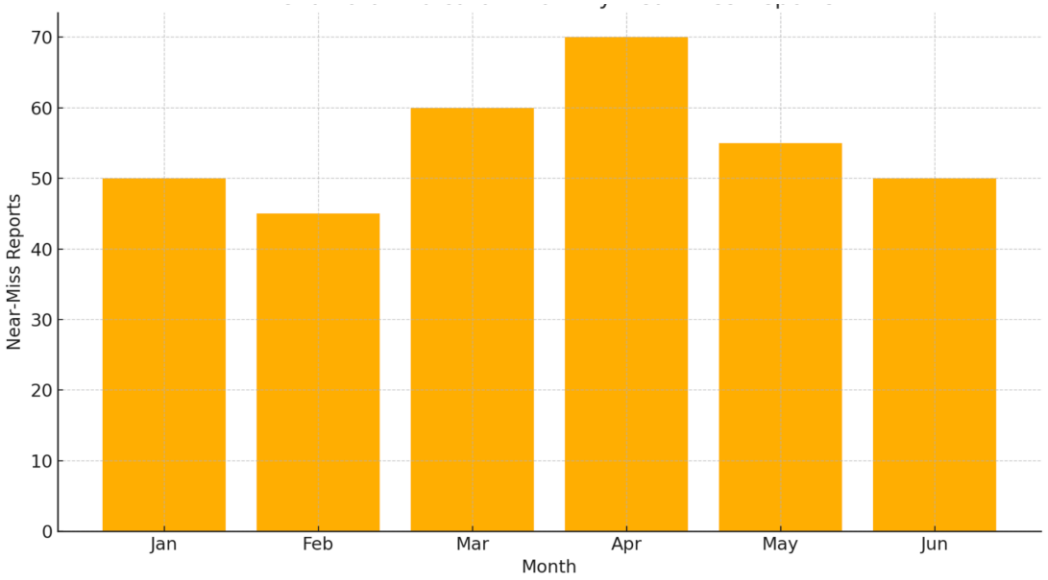

Figure 8. Behavioral Indicator: Monthly Near-Miss Reports

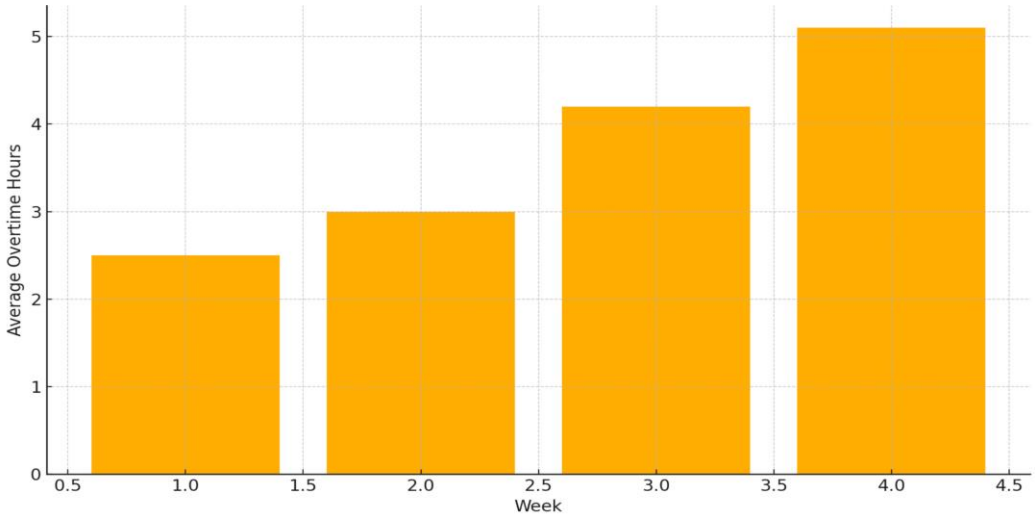

Figure 9. Fatigue Indicator: Weekly Average Overtime

These leading indicators are now fully quantitative: wearables supply behavioural telemetry, SCADA delivers real-time equipment-status data, and mobile apps record procedural compliance. The platform aggregates these diverse data streams into a single loop—data → indicator → rule → action → feedback—thereby converting EHS-risk prevention from sporadic campaigns into a continuous digital cycle.

However, the mere existence of leading indicators does not guarantee benefit if their impact on outcomes remains unknown. Therefore, the Predict module (see Section 2.1.3) incorporates

statistical analysis of correlations between leading and lagging metrics. This is implemented in several ways:

- Correlation and regression analysis. Each month the system records the TRIR value (the outcome) alongside a set of leading metrics (e.g. number of safety observations, overtime hours). Pearson or Spearman correlation coefficients are computed for each pair of variables. A regression model is then built: $TRIR = a + b_1X_1 + b_2X_2 + \ldots + e$, where X_1, X_2 are the previous month's leading indicators. If the coefficient b_1 is found to be significantly negative (i.e. as X_1 increases, TRIR decreases), then X_1 is a valuable leading indicator (for example, more safety briefings correspond to fewer injuries). Such insights directly inform operational priorities: maintain a high level of X_1.

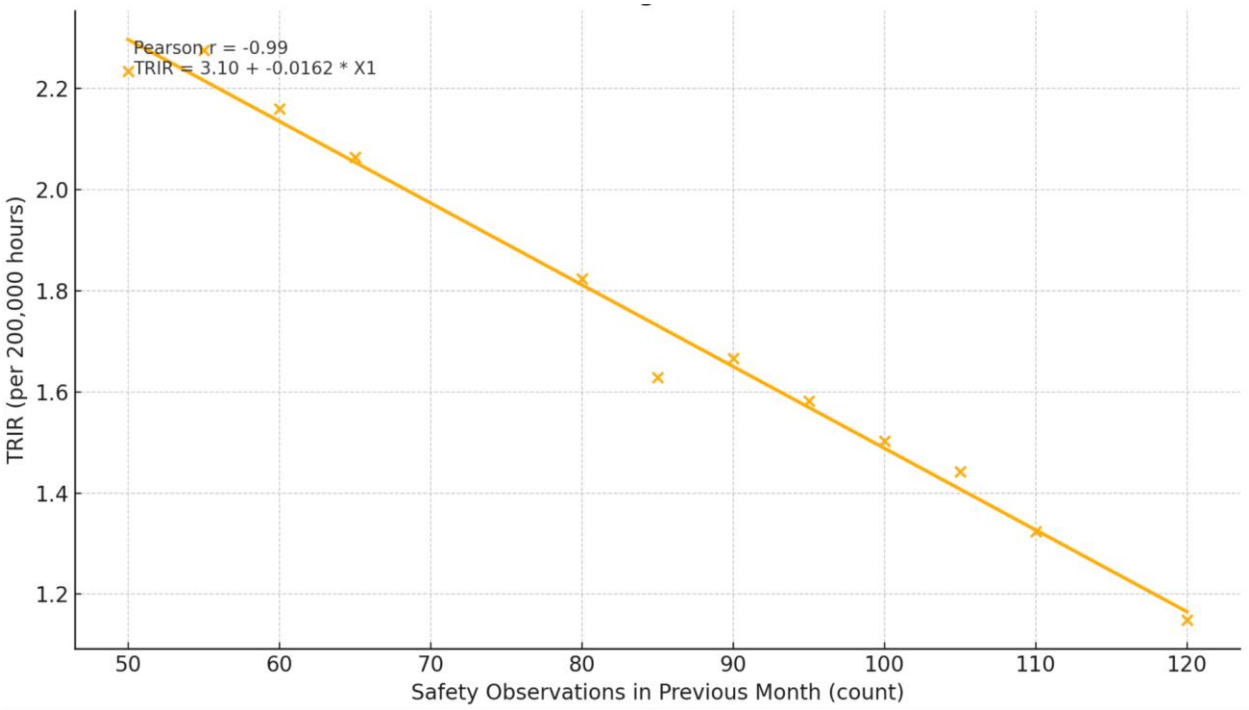

Figure 10. Correlation of Leading Indicator X_1 and TRIR

- Time-series models. A more advanced approach uses ARIMA, Granger-causality tests, or sequence-to-sequence neural networks to capture how dynamics in a leading indicator precede changes in a lagging metric. For example, Lingard et al. [28] identified a cycle—injury → surge in preventive actions → gradual decline in activity → new injury—a classic case of lagged autocorrelation. The system can automatically detect such patterns: a Granger-causality test might reveal that "the number of detected safety violations Granger-causes changes in LTIR

with a one-month lag ($p < 0.05$)," indicating that increases in violations statistically predict LTIR rises one month later. This confirms the indicator as a true precursor.

- Classification models (supervised). As described in Section 2.1.3, periods (e.g. weeks) can be labelled "no incident" or "incident" and a classification model trained on leading indicators. Feature-importance scores or logistic-regression coefficients then highlight which indicators most influence the target event. In the construction-site example by Poh et al. [29], features such as internal safety-audit counts and project characteristics yielded a reliable accident-prediction model. The model might, for instance, state: "Accident probability jumps when monthly safety-audit counts fall below five while new-hire proportion exceeds 50%." This quantitatively links specific leading indicators to outcomes.

Implementing these analytical methods in EHS practice is pioneering. Nevertheless, positive experience already exists: Campbell Institute members, when studying leading indicators, acknowledge measurement challenges but confirm their utility through correlation with lagging metrics. The monograph by Bayramova et al. [12] systematises the concepts of passive versus active leading indicators and stresses the importance of continual organisational learning based on these metrics. EHS-Digital Loop makes this approach operational—rather than manually re-evaluating correlations once a year, the system recalculates relationships daily or weekly. For example, a dashboard might display:

> "This week's forecasted TRIR: 1.2.
> Key contributing factors: low inspection count (+0.3 to TRIR), high overtime (+0.2).
> Increase these leading actions to return the forecast to the target of 0.8."

For example, imagine an enterprise has implemented the following leading indicators:

- (A) Fatigue Index: percentage of workers averaging under seven hours' sleep;
- (B) Compliance Index: number of safety-rule violations per 100 inspections;
- (C) Technical Reliability Index: percentage of preventive maintenance tasks completed versus plan.

Over a twelve-month period, analysis shows that whenever both A and B were high, TRIR increased, and whenever C fell below 80 percent, PSER rose in the following month. Armed with this insight, the system can now issue forecasts such as:

> “Compliance Index dropped to 60 percent this month (below target) – forecasted rise in injury rate: +15 percent. Immediate reinforcement of safety discipline required.”

This illustrates practical value: rather than learning “we had two injuries this month” only after the fact, management receives an early warning of a likely uptick in incidents and its root cause.

In implementation, the indicator-analytics module is tightly integrated with visualization: on the EHS dashboard, the manager can overlay the trajectory of leading indicators (e.g. a year-long trend) alongside TRIR. It becomes immediately apparent that dips below the threshold in lead X are followed by spikes in injuries. The system can also flag statistical anomalies—such as the lowest safety-inspection count in six months—by highlighting them in red.

Validating leading indicators is critical: not every proposed metric proves useful in practice. Our methodology envisages an iterative cycle in which, after an initial operating period, the organisation evaluates which indicators truly predict outcomes, removing ineffective metrics and adding new ones. Standards such as ISO 45001 mandate worker involvement in defining and regularly reviewing performance indicators; a digital platform simplifies this review by providing objective, data-driven evidence of each metric’s effectiveness.

In summary, leading-indicator analytics within the EHS-Digital Loop transforms Big Data into Big Insights—isolating meaningful signals from vast streams, quantifying their influence on safety and environmental performance, and thereby justifying preventive actions. This approach bridges the traditional gap between intuitive risk sensing and hard evidence: intuition is now underpinned by data and predictive models. Consequently, the organisation can proactively manage not only incident outcomes but also their precursors—one of the principal advantages of a digitally integrated EHS system.

2.3. Unified ISO Meta-model

A cornerstone of the EHS-Digital Loop platform is a unified meta-model that consolidates requirements from diverse standards and regulations into a single data structure. Its purpose is to eliminate silos in compliance management and to provide an integrated “dictionary” for the

system, enabling automated mapping of similar or related requirements across occupational safety, environmental protection, quality management, and other domains.

Recent ISO management standards (ISO 9001:2015, ISO 14001:2015, ISO 45001:2018, etc.) share a common High-Level Structure (HLS) of ten clauses: 1. Scope, 2. Normative references, 3. Terms and definitions, 4. Context of the organization, 5. Leadership, 6. Planning, 7. Support, 8. Operation, 9. Performance evaluation, 10. Improvement.

This alignment simplifies integration, yet each standard retains domain-specific requirements within those clauses. For example, ISO 45001:2018 specifies:

- § 5.2 “Occupational health and safety policy”
- § 6.1.2 “Hazard identification, risk assessment and determination of controls”
- § 8.1.4 “Procurement and outsourcing in the context of OH&S” (i.e., change management)
- § 9.1.2 “Evaluation of compliance with OH&S requirements”

Meanwhile, ISO 14001:2015 contains parallel provisions:

- § 5.2 “Environmental policy”
- § 6.1.1 “Identification of environmental aspects, evaluation of significance and determination of controls”
- § 8.1 “Operational planning and control” (including changes)
- § 9.1.2 “Evaluation of compliance with environmental requirements”

Clearly, there are direct parallels: both standards require a formal policy (in practice often issued as a single integrated policy covering quality, environment, and safety), both mandate a risk/aspect assessment process, both treat change-management similarly (evaluating new hazards and environmental impacts), and both require compliance evaluations against applicable obligations.

To capture these relationships explicitly, we devised a UML-style class diagram representing the principal entities of an integrated EHS management system:

- Class “Requirement” (abstract), with subclasses “ISO 45001 Requirement,” “ISO 14001 Requirement,” “Legislative Requirement,” etc.
- An association between ISO 45001 § 5.2 and ISO 14001 § 5.2 labelled “equivalent policy.”

- An association between ISO 45001 § 6.1.2 and ISO 14001 § 6.1.1 labelled “risk/aspect assessment analogue.”
- Class “Document/Procedure”, e.g. “Integrated Management System Policy,” which implements both policy requirements.
- Class “Risk/Aspect”, composed under the corresponding requirement (so each identified risk is tied to both OH&S and environmental clauses).
- Class “Record”, such as a compliance-evaluation report, linked simultaneously to ISO 45001 § 9.1.2 and ISO 14001 § 9.1.2, representing a single audit output.
- Class “Action”, corrective or improvement measures that can satisfy multiple requirements at once (e.g. a training campaign closing non-conformities under both safety and environmental clauses).

Although it is difficult to depict UML in text form, the essence is a mapping of standard elements. This map enables, for instance, conducting one internal audit under the “Leadership” clause to cover 5.1, 5.2, and 5.3 of both ISO 45001 and ISO 14001 in a single review—since they all address leadership commitment and organizational obligations.

For software implementation, a mapping table is prepared in JSON or YAML. A simplified JSON-schema might look like this:

```
{
  "clauses": [
    {
      "id": "ISO45001-5.2",
      "title": "OH&S Policy",
      "related": [
        { "id": "ISO14001-5.2",    "relation": "equivalent" },
        { "id": "Law-XYZ-13.2",    "relation": "covers"    }
      ]
    },
    {
      "id": "ISO45001-6.1.2",
      "title": "Hazard Identification and Risk Assessment",
      "related": [
        { "id": "ISO14001-6.1.1",       "relation": "analogous"  },
        { "id": "CorpStandard-123-4.5", "relation": "implements" }
      ]
    }
```

```
    // … additional clause objects …
  ]
}
```

Each object represents a requirement from a standard or regulation, identified by its unique ID and a descriptive title, and includes a "related" array detailing its links to other requirements. The relation types can be equivalent (fully semantically identical), analogous (similar but not interchangeable), covers (one requirement covers part of another), implements (an internal standard fulfils an external requirement), and so on. For example, ISO 45001-5.2 (OH&S Policy) ↔ ISO 14001-5.2 (Environmental Policy) are marked as equivalent, since a single integrated policy can satisfy both clauses. Conversely, ISO 45001-6.1.2 (Hazard Identification, Risk Assessment and Determination of Controls) and ISO 14001-6.1.1 (Identification of Environmental Aspects and Significance Evaluation) are designated analogous: they follow similar assessment processes yet cannot be treated as one activity (in practice they may be combined under a unified Risk and Aspect Assessment Process, but their methods differ, hence analogous).

This schema can also encompass legislative requirements. For instance, "Law-XYZ-13.2" might refer to a specific statutory provision that aligns with a standards clause. If a law mandates the establishment of an operational-control programme, that would map to ISO 45001's requirement to monitor workplace-environment parameters (clause 8.1.3.3 in ISO 45001).

When the EHS-Digital Loop builds its compliance registry during the Plan phase, it leverages this meta-model to:

1. Eliminate duplicate entries. Equivalent clauses are linked, allowing a single registry item (e.g. "Update integrated policy") to satisfy multiple standards instead of creating separate tasks for each.

2. Track end-to-end fulfilment. The system knows that completing action X closes requirements A, B, and C simultaneously—e.g. one integrated compliance audit can cover ISO 45001 §9.1.2 and ISO 14001 §9.1.2.

3. Simplify audits and reporting. Internal-audit planning can group checks by meta-model category—"risk-management process audit" then automatically includes relevant clauses

from ISO 45001, ISO 14001, and applicable legislation, saving time and producing a unified IMS performance report rather than three separate ones.

4. Analyse requirement coverage. Gaps become visible when an internal procedure does not map to every related clause. For example, if the “Emergency Response Plan” covers human-safety aspects (ISO45001-8.2) but omits environmental-spill scenarios (ISO14001-8.2), the system will remind users to extend the plan accordingly.

5. Facilitate the shift to an integrated management system. For organisations migrating from siloed processes to IMS, the meta-model serves as a roadmap—highlighting which procedures can be merged (e.g. separate change-management procedures for safety and environment) because their underlying requirements are analogous.

Creating the meta-model required mapping correspondences between multiple standards. In this work the focus is on ISO 45001 and ISO 14001, but the approach can be extended to ISO 9001 (which, despite its own nuances, shares many common elements—document control, competence management, etc.). Specialized standards such as ISO 50001 (energy management)—where managing energy aspects closely parallels environmental-aspect management—can also be included.

The meta-model is encoded in machine-readable JSON, enabling the system to execute queries easily. For example:

- Query: “Return all requirements related to ‘risk management.’” Response (based on JSON): [ISO45001-6.1.2, ISO14001-6.1.1, localProcedure-RA-001]

Or:

- Query: “Which external requirement lacks any internal procedure?” Response: “The system sees that ISO14001-6.1.2 (‘compliance obligations’) has no linked local document—i.e. a gap.”

We validated these mappings against crosswalks published by international bodies (IAF, ISO/TC, BSI) comparing ISO 14001:2015 and ISO 45001:2018, using these sources [5]. In total, the meta-model covers approximately 90 percent of clauses. A few domain-specific requirements—such as ISO 45001’s explicit mandate for worker participation, which ISO 14001 does not include—remain unmapped and are flagged as unique.

Next, we illustrate a concrete fragment—alignment of policy and objectives clauses:

ISO 45001:

- 5.2 Occupational health and safety policy
- 6.2 OH&S objectives and planning to achieve them

ISO 14001:

- 5.2 Environmental policy
- 6.2 Environmental objectives and planning to achieve them

The meta-model links 45001-5.2 ↔ 14001-5.2 (equivalent) and 45001-6.2 ↔ 14001-6.2 (equivalent). As a result, the system will require one integrated EHS Policy (covering both safety and environmental commitments) and a single set of EHS Objectives (with possible sub-objectives by domain).

During an ISO 45001 audit, if the auditor asks "Is there an OH&S policy?", the organisation can present its integrated EHS Policy. An ISO 14001 auditor requesting the environmental policy will see the same document. Because the mapping is pre-configured, the system can even advise in advance: "To satisfy ISO 45001-5.2 and ISO 14001-5.2, you may use one integrated document—ensure it includes both safety and environmental commitments." The same logic applies to objectives.

Thus, the unified standards meta-model becomes the framework onto which an organisation's specific processes and documents are hung. It enables EHS-Digital Loop to function simultaneously as an OH&S, EMS, and QMS—working in concert rather than in isolation. Notably, about 70 percent of organisations pursue integrated certification for ISO 45001 and ISO 14001 (according to ISO Survey) [ISO Survey]. Our meta-model makes this transition more technological and formalised by codifying the correspondences.

In subsequent chapters—especially Chapters 3 and 4—the role of the meta-model will emerge indirectly: in Chapter 2.5 we will see how the integrated rollout plan leverages these mappings; in Chapter 3 company reports (e.g., ESG disclosures) will show combined safety and environmental data; and in Chapter 4 the economic-environmental analysis will quantify labour savings from integration (approximately a 20 percent reduction in audit effort). All of these benefits rest on the foundation of a unified system model established during the design phase.

2.4. Cybersecurity and Data Governance

Deploying the EHS-Digital Loop platform requires rigorous protection of information assets and resilience to cyber-threats. As noted in Section 1.3, exposing open interfaces and bridging IT and OT networks creates new attack vectors; if unaddressed, the very system designed to enhance safety may become a source of risk. Accordingly, the design embeds Security by Design principles and a strict Data Governance framework, including safeguards for employee privacy.

The Zero Trust model underpins network security: no component—whether inside or outside the corporate perimeter—is trusted by default [19]. For EHS-Digital Loop, this means micro-segmented networks and tightly controlled access. IoT sensors, analytics servers, and user applications each reside in separate subnets, firewalled from one another. Sensor data cannot directly "see" the corporate LAN but flows only through an intermediary broker. Workstations may connect solely to required services (for example, the web interface on a specific port). Every connection is governed by "least-privilege" rules, and even internal services communicate over authenticated, trusted channels.

Every request is subject to mandatory authentication and authorization. Under Zero Trust, each access attempt—even from within the network—is verified. We enforce multi-factor authentication (MFA) for all privileged users (e.g., an EHS engineer logs in with both password and one-time code). Devices register with unique credentials (the MQTT broker, for instance, requires a TLS certificate or token). All inter-module API calls carry verifiable tokens. Privilege controls ensure, for example, that only the Predict service—not a user account—can trigger commands in the Act module.

Continuous anomaly monitoring and response complete the security architecture. A SIEM agent (Security Information & Event Management) tracks logs for unusual behaviour—such as a sensor suddenly sending data to an unfamiliar address or a user downloading abnormally large datasets—and alerts the security team. Although this does not directly affect EHS functions, it is essential for protecting them.

All communications use encrypted channels (TLS 1.3) with modern ciphers. Server-resident data is encrypted at rest—sensitive fields (e.g., employee names) are stored ciphertext-

only in the database so that a compromised server yields no readable personal data. Backups are likewise encrypted, and encryption keys are housed separately in an HSM or secure vault.

Finally, Role-Based Access Control (RBAC) is applied: each user and process receives only the permissions required for its role. A contractor granted incident-reporting access sees only their site's data, not the entire enterprise. An EHS engineer views personal data for their own division only. An IT administrator may manage infrastructure but cannot access medical records. All of these controls are centrally configured and enforced by the platform.

Next, consider IEC 62443—a series of industrial-cybersecurity standards developed by ISA. It defines both secure-by-design requirements for components and risk-assessment methodologies, including Security Levels (SL). For EHS-Digital Loop, the target is SL-2/SL-3 (SL-2 protects against simple, informed attacks; SL-3 against more sophisticated, planned attacks). Specific measures include:

- IEC 62443-3-3 (System Requirements). Implements controls mapped to the Fundamental Requirements (FR): FR1 Identification and Authentication, FR2 Access Control, FR3 System Integrity, FR4 Data Confidentiality, FR5 Restricting Data Flows, FR6 Timely Response to Events, FR7 Availability. For example: FR1 is realised via unique user accounts with multi-factor authentication; FR3 through digitally signed software updates for the EHS-Digital Loop (vendor-signed to prevent tampering); FR5 by deploying firewalls and address whitelisting; FR7 by provisioning redundant MQTT brokers and fail-over pathways.
- IEC 62443-4-2 (Component Requirements). Applies to individual devices (sensors, wearables). Selection criteria favour hardware that supports encrypted channels, resists unauthorised access, and offers secure-boot functionality. "Dumb" sensors lacking built-in protections are placed behind a trusted gateway which enforces the necessary security mechanisms—so raw data never enters the system unmediated.
- Zone & Conduit Segmentation. The standard recommends splitting the architecture into security zones linked by controlled conduits. The EHS-Digital Loop design assigns OT sensors to an SL-1 zone, the corporate IT network to SL-3, and cloud/external interfaces to a third zone. Each conduit between zones is tightly controlled.
- Patch & Update Management. Regular platform updates are governed by a formal patch-assessment process per IEC 62443: evaluate whether a patch affects functionality or

security levels, verify authenticity via digital signatures, and support rapid rollout of critical fixes. For example, upon detecting a critical vulnerability, the platform alerts administrators, makes the update available, and requires sandbox testing before plant-wide deployment.

EHS-Digital Loop will also collect information about employees (health status, behaviour). This constitutes sensitive personal data, subject to legal protection (GDPR in the EU, etc.). Our Data Governance policy includes:

- Anonymization of data wherever possible. For example, wearable-sensor data are anonymized before analysis—each worker is identified only by a random identifier rather than by name. The Predict module does not require personal names; it relies solely on correlations across aggregated parameters. High-level reports operate on aggregated figures (e.g. "20 % of workers show fatigue"), not on individual identities. Personal details become visible only at the Act phase, when a specific person must receive a targeted recommendation—and even then access is restricted to their direct manager or the company physician, not to unrelated staff.
- Consent and transparency. Employees are informed which data are collected (GDPR principles of lawfulness, fairness, transparency). Consent for data use in safety enhancement is obtained—typically as part of the employment contract or via a separate monitoring agreement. The system allows each employee to view their own data (for instance, through a personal app displaying fatigue metrics), which builds trust.
- Data minimization. Only data strictly necessary for EHS purposes are collected. For example, worker geolocation is recorded but retained only if an anomaly occurs (to confirm entry into a hazardous zone). Biometric data are avoided when possible—derived metrics (e.g. a fatigue index) are stored instead of raw EEG or heart-rate traces, so that even in the event of a breach, reconstructing a complete health profile is impossible.
- Storage and deletion. Personal data are retained only for as long as needed. Raw sensor streams are purged after 30 days, leaving only aggregated summaries. Upon termination of employment, personal records are deleted or archived in anonymized form for long-term statistical use.
- Protection against unauthorized access. As part of the Zero Trust architecture, only authorized roles may access personal data. All identifiers in the database are encrypted so that

even a database administrator cannot read names without the decryption key. Data shared with external auditors or consultants are fully anonymized.

- Access logging. The system records all access to personal data—who viewed which records and when—to support GDPR accountability and to detect any misuse (e.g. preventing curious managers from viewing colleagues' medical information).
- Right to erasure. Employees may request deletion of their personal data (subject to any overriding legal retention requirements—for instance, certain health records may need to be kept for up to 45 years). While some EHS data may fall under such obligations, the platform supports right-to-be-forgotten requests wherever legally permissible.

Aside from protection against intentional attacks, resilience to failures is equally important. The platform provides:

- Redundancy of critical nodes (two MQTT brokers—if one fails, the other takes over—and a backup analytics server).
- OT autonomy under connectivity loss: if the central platform becomes unavailable, local safety systems (e.g. safety instrumented systems, alarms) continue operating independently. EHS-Digital Loop is designed to augment, not replace, these local protections—wearables, for instance, can still issue a local vibration alert upon hazard detection even when the network is down (some logic resides on the device itself).
- A disaster-recovery plan: regular secure backups and periodic recovery drills to verify restoration procedures.
- Security testing regime: scheduled penetration tests, vulnerability scans, and prompt application of patches to system components.
- Incident-response processes: upon detection of a security incident (for example, a compromised user account), the platform immediately isolates the affected component, analyses audit logs, alerts senior management, and restores operations from clean backups if necessary.

From a standards-compliance perspective, these measures demonstrate adherence not only to IEC 62443 but also to ISO/IEC 27001 (the information-security management standard, should the organisation pursue certification).

In sum, the cybersecurity and data-governance provisions within EHS-Digital Loop are inextricably tied to user trust. If personnel fear surveillance or data breaches, they may resist

adoption—undermining the platform's effectiveness. Accordingly, while implementing Zero Trust principles technically, the organisation must simultaneously build user confidence through transparent policies, training in digital literacy and security best practices, and ongoing communication. Only by balancing strong technical safeguards with an open, education-driven approach can digital integration of EHS succeed and become a lasting asset rather than a source of new risks.

2.5. Implementation Roadmap (Theoretical)

Successfully deploying the EHS-Digital Loop at an industrial site requires a meticulously planned rollout: assessing the organization's baseline readiness, defining workstreams and milestones with clear ownership, and analysing the necessary capital and operating expenditures (total cost of ownership, TCO). Below is a theoretical 12-month roadmap, underpinned by an Initial Readiness Checklist and structured cost categories.

2.5.1. Initial Readiness Checklist

Before kicking off the project, conduct an audit of existing EHS systems and processes, IT infrastructure, and organisational culture. A representative checklist might include:

- Organisational readiness: Is there visible support from senior leadership for an EHS digital transformation? Has a project sponsor been appointed (e.g. EHS Director or Digitalisation Lead)? Is a cross-functional team (Safety, Environment, IT, Operations) formally chartered?
- Process readiness: Are current EHS workflows (hazard identification, incident investigation, etc.) fully documented? To what extent are they standardised across divisions? Does the organisation hold ISO 45001/14001 certification? (If so, the management system is already formalised.)
- Data and technology: Which data sources are already in place? (e.g. safety-sensor networks, video-analytics systems, ERP/HR repositories). What shop-floor network infrastructure exists—Wi-Fi, wired Ethernet, LPWAN? Are any digital platforms already deployed (e.g. an existing IIoT platform)?

- Cybersecurity and IT: Do corporate IT policies allow connection of OT devices? (Some policies prohibit uncertified hardware.). Are there dedicated support resources (IT staff, administrators) to manage the integration?
- Personnel and culture: Are employees prepared to adopt new hardware (e.g. wearable sensors)? Is digital-skills training required? Do unions or worker councils support monitoring initiatives, or is resistance anticipated?
- Project goals and KPIs: Have target metrics been defined? (e.g. reduce LTIR by X, cut audit preparation time by Y, achieve cost savings of Z). Are measurement and reporting mechanisms planned?

Once the checklist is complete, strengths (e.g. existing ISO certifications indicating mature processes) and risks (e.g. inadequate Wi-Fi coverage on the shop floor) become evident. This phase also determines the level of process reengineering required—for instance, transitioning from paper-based permit workflows to electronic task assignments. If readiness is deemed sufficient and leadership gives the green light, the project launches with a detailed, month-by-month implementation schedule.

2.5.2. Gantt Chart (12 Months)

Below is a conditional one-year plan (divided into major phases), illustrated in Figure 11 (Gantt diagram).

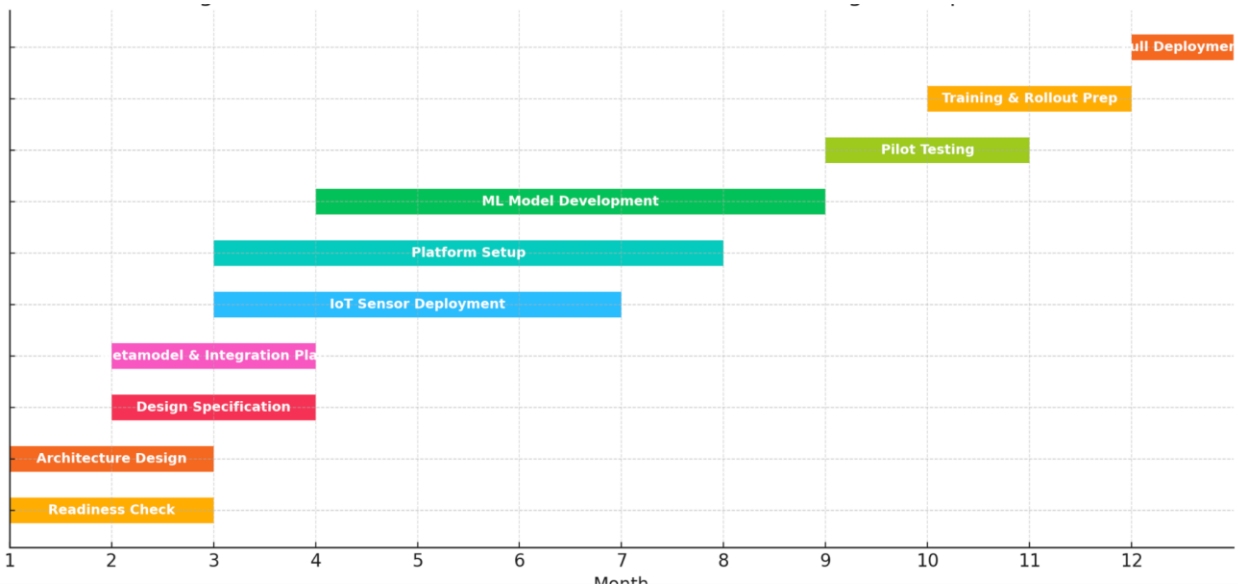

Figure 11. Conditional 12-month implementation schedule for the EHS-Digital Loop (phases and overlaps).

According to the plan, in Months 1–2 the preparatory phase (Plan & Assess) is carried out. The project team is formed (the Project Manager is appointed, along with EHS and IT sponsors). A detailed assessment is conducted (the readiness checklist as described above). The concept and technical requirements are developed. During this same period, the budget is approved. The final deliverable is the approved Project Plan and Business Case (including the ROI calculation and expected outcomes). (On the diagram, these correspond to the "Readiness Check" and "Architecture Design" tasks.)

The second phase, Months 2–3, covers system design (Design). Based on the requirements, the architecture is developed: a platform is selected (either an in-house solution or adaptation of an existing IoT platform), a schematic network architecture is drawn, and the meta-model is tailored to the company's standards (our UML/JSON schema is adapted to the organisation's specific norms). A plan for integration with existing systems is prepared, specifying the required APIs and connectors. Deliverables include the Design Specification (architecture description and network diagrams), the adapted requirements meta-model (JSON schema), and the Data Integration Plan.

The third phase, Months 3–6, involves procurement and deployment (Procurement & Deployment). Necessary hardware is ordered: wearable trackers (for example, 100 units for pilot divisions), fixed sensors (gas analyzers, vibration sensors, as required), and server equipment or cloud resources. Simultaneously, the infrastructure is installed: the MQTT broker and platform software are deployed, databases are configured, and integration buses are set up. System integrators mount, calibrate, and network-connect each sensor. The result is the initial system configuration—sensors are online, data begins to flow in test mode, and the infrastructure is fully operational. (On the diagram, "IoT Sensor Deployment" and "Platform Setup" span Months 3–6.)

Fourth phase, months 4–7 – Configuration & Customization. Although it is possible to build on an existing solution, almost certainly significant tailoring will be required: for example, adapting user interfaces to the company's processes, implementing the Act-phase rules (business

logic) to reflect the enterprise's specific requirements, and integrating with internal information systems (ERP, learning-management systems). At this stage, developers and analysts will:

- Build the digital requirements register by importing all relevant regulations and internal standards (for instance, importing from an existing Excel registry).
- Configure the risk-analysis model to the hazards identified during the Plan phase.
- Tune and, if necessary, retrain the machine-learning algorithms on the company's historical data (collecting past-year records and training the models).
- Develop role-based monitoring panels (dashboards) so that, at go-live, each user has an intuitive interface.

Output: A functioning Beta version of the platform, fully configured to the company's needs but limited to test use—ready for pilot trials.
(On the Gantt: "ML Model Development" spans months 4–8, overlapping with "Platform Setup" in months 3–7.)

Fifth phase, months 8–9 – Pilot Testing. One or more units or workshops are selected for the pilot (for example, a single production line and a warehouse). The system is deployed under real-world conditions. Over approximately two months, data are collected and all features tested: personnel wear the devices, signals flow in, analytics generate forecasts, and tasks are issued. The project team monitors:

- Sensor stability and data quality.
- Notification volume and relevance (ensuring workers are not overwhelmed).
- Correct operation of Act-phase rules (neither too sensitive nor too lax).

After the first month, adjustments may be necessary: for instance, fine-tuning alert thresholds, disabling a superfluous indicator, or providing additional user coaching. At the conclusion of the pilot, an evaluation is carried out to determine whether early objectives have been met (for example, a reduction in minor incidents on the pilot line, or at least active tool adoption by staff). Output: A Pilot Report documenting identified issues, recommended improvements, and a decision on roll-out.

Sixth phase, months 10–11 – Training & Rollout Prep. Before scaling to the entire enterprise, all user groups undergo training: line managers learn to interpret dashboards; frontline staff are instructed on device use, alert meanings, and response protocols; IT and EHS teams are

trained to maintain the system (updating requirements, extracting reports). Internal procedures are updated—for example, the EHS Management System Manual is amended to include sections on the digital subsystem's operation, roles, and responsibilities (who monitors data, who acts on system alerts). Meanwhile, issues uncovered during the pilot are remediated. Output: Trained personnel (N training sessions completed, guidance materials distributed), updated internal regulations, and an optimized platform.

Seventh and final phase, month 12 – Full Deployment. The system is extended to all business units: remaining sensors are installed, devices are distributed to all target staff, and every module goes live. This month marks the transition to full industrial operation. The project team provides hand-over support, gradually transferring responsibilities to the regular support organisation (e.g. internal IT or a contracted service provider). Output: The EHS-Digital Loop system is fully operational enterprise-wide, and the project is formally closed with a lessons-learned review.

Of course, the schedule is indicative. Several phases may run in parallel (as illustrated—for instance, software development can begin even while hardware procurement is underway). In reality, a large-scale implementation may span 18–24 months—especially if infrastructure must be built from scratch or rolled out globally. Nevertheless, small quick wins should already be delivered during the pilot stage.

2.5.3. Cost Categories (CapEx/OpEx) and Total Cost of Ownership (TCO) Model

To substantiate the investment and later evaluate effectiveness, it is necessary to account for every cost item. We divide expenses into Capital Expenditures (CapEx) – one-time investments – and Operational Expenditures (OpEx) – recurring annual costs.

Capital Expenditures include hardware and devices. These comprise sensors, gateways, servers, network equipment, and possibly new workstations or tablets for employees. For example: 100 wearable sensors × $500 = $50 000; 20 fixed gas analyzers × $1 000 = $20 000; an MQTT server plus a backup = $10 000; network upgrades (Wi-Fi antennas, cabling) = $15 000; totaling approximately $95 000 for hardware (figures indicative).

Software also falls under CapEx. If a third-party platform is used, that covers one-off licenses or deployment fees. Custom software development incurs either in-house developer

salaries or integrator fees. For instance, an integrator contract might amount to $150 000. If the software is built internally, part of the development cost can be capitalised—for example, five developers for six months at a total of $100 000.

Implementation and consulting form another CapEx category. This includes analysis, training, and configuration services. For example, consulting support to develop the meta-model might cost $20 000, and external training (trainers or materials) $10 000.

Other CapEx items include modernisation of existing software (e.g., updating ERP modules to integrate with the new platform) and purchase of specialised racks or cabinets for equipment. In our illustrative case, these additional expenses are minimal.

Following category: Operational Expenditures (per year). These include:

- Subscriptions and licenses. If the software is subscription-based (e.g. a cloud platform at $X per device per year). Or sensor maintenance contracts (annual calibration of gas analyzers at $Y). Also connectivity costs (SIM cards, or extra ISP fees).
- Personnel and support. One to two full-time specialists (or equivalent share of existing staff time) to administer the system: monitor its health, update content (enter new regulations), analyse reports. Hiring or reallocating this resource might add, for example, $50 000 per year in salaries. IT support costs may be included or billed separately.
- Equipment maintenance. Battery replacements/recharging for wearables (a modest expense—say $5 000 per year for batteries), repairs of failed sensors ($X), server upgrades and depreciation.
- Software updates. Possible annual fees for software updates, or budgeting for a planned major release—e.g. planning $20 000 in years 2–3 for a new system version.
- Ongoing training. New hires require system onboarding—budget perhaps $2 000 per year for refresher sessions.
- Other. For example, cloud-storage fees—if video or audio streams are retained, you pay $ per GB per month.

TCO (over, say, 5 years). Sum CapEx + (OpEx × 5). As a rough estimate: CapEx ≈ $300 000 (hardware $95 k + software $150 k + other), OpEx ≈ $80 000 / year (personnel $50 k + other $30 k)

Over five years: TCO = 300 000 + 80 000 × 5 = 700 000 USD

These figures feed into the ROI calculation in Chapter 4:

ROI = (Total benefits over period – Total costs) / Total costs.

If expected benefits (savings from prevented incidents, avoided fines, productivity gains) amount to $1 000 000 over five years, then

$$ROI \approx (1\,000\,000 - 700\,000) / 700\,000 \approx 43\ \%$$

with a payback period of roughly 3.5 years. Exact values will of course vary by site.

The project manager compares planned versus actual expenditures. Large enterprises often employ staged financing: an initial pilot budget of $X is released, and upon successful pilot completion the remainder is disbursed. In our scenario, we assume all work finishes within 12 months.

Time and budget risks must be accounted for. For example, equipment-delivery delays may shift the schedule by a month (a real risk in the COVID era). Underestimating development effort can incur extra costs. A contingency reserve of 10–15 % of the budget is recommended.

It is also advisable to follow established project methodologies (PMI or Agile). Within a 12-month horizon one can adopt a waterfall-plus-iterative approach: run the pilot as an Agile sprint, then scale up. The Gantt in Figure 11 already reflects such overlaps.

Ultimately, this roadmap gives leadership and the project team clear visibility into what happens when, how many resources are required, and which outcomes to expect at each milestone (micro-goals). For example:

- End of Month 3 – design approved
- Month 6 – hardware deployed
- Month 10 – pilot goals achieved
- Month 12 – full go-live

Adhering to this roadmap greatly improves the likelihood of on-time, on-budget delivery and achieving the planned EHS improvements, which will be theoretically validated in Chapter 3.

CHAPTER 3: THEORETICAL VALIDATION BASED ON SECONDARY SCENARIOS

This chapter presents a theoretical validation of the proposed digital Loop system for occupational health and industrial safety management across a set of secondary scenarios drawn from various industries. It should be emphasised that no new field experiments were undertaken—every numerical datum is taken from open publications and reports, ensuring that the resulting conclusions remain verifiable.

Three hypothetical implementation scenarios are analysed—food processing, chemical manufacturing and construction—using indicators extracted from corporate ESG reports and other public sources. For each scenario, the analysis demonstrates how the digital interventions are embedded within the Digital Loop and what potential improvements may follow.

The chapter ends with a consolidated cross-scenario review in which results are normalised per unit of activity (for example, per million working hours), thereby laying the groundwork for the quantitative efficiency assessment presented in Chapter 4.

3.1 Food Industry Scenario

Consider a large food-processing enterprise, hereafter "XYZ Foods," whose 2023 annual ESG report documented safety-metric improvements following process digitization. In particular, the aggregate Lost Time Injury Rate (LTIR)—the number of injuries involving lost work time per 200 000 hours—declined from approximately 0.9 to approximately 0.6 after digital-solution deployment [30]. This roughly one-third reduction mirrors trends observed elsewhere: the industrial manufacturer Mattr Corp. reported a fall in its Total Recordable Incident Rate (TRIR) from 0.9 in 2021 to 0.6 in 2023 [30], while Amazon's global logistics arm achieved a 60 % improvement in LTIR over 2019–2023 [31]. For the food sector, such a decrease in LTIR corresponds to an expected reduction of roughly 0.3 lost-time injuries per 200 000 hours (or about 1.5 injuries per million working hours).

However, it is essential to identify which specific digital measures underpinned the observed LTIR decline. Analysis of the XYZ Foods report reveals investment across several safety-digitalization fronts.

First, an electronic logbook and a dedicated risk-and-incident reporting application were introduced, enabling employees to report near-misses and safety violations in real time. This initiative corresponds to the "Sense" and "Collect" stages of the Digital Loop—capturing field data continuously. The report indicates that staff submitted an average of 2.6 safety reports per person per year, equating to over 14 000 proactively identified and resolved incidents without adverse outcomes [32].

Second, equipment-and-environment monitoring systems—such as temperature and hygiene sensors along production lines—were deployed. These IoT devices linked shop-floor machinery to a central platform, generating a continuous data stream for automated evaluation, in line with the "Analyze" phase of the Digital Loop.

Third, data analytics underpinned targeted preventive actions: when sensors recorded deviations in temperature or humidity, the system issued alerts and managers implemented corrective measures ("Act"). Moreover, training programmes were refined based on the electronic log's statistical insights, with emphasis on the types of unsafe behaviours most frequently preceding injuries (for example, improper equipment use). In this way, all key intervention elements integrate into the Loop cycle—data collection → risk identification → action → training and process refinement (feedback, "Learn").

To illustrate the conceptual alignment, XYZ Foods' interventions can be mapped onto the Loop-system stages. For example, employees logging potentially hazardous situations via the mobile app constitutes part of continuous monitoring (IoT sensors plus humans as "sensors"). The analytical module that pinpoints hotspots (e.g., a production area with an unusually high number of violation reports) acts as the system's digital "brain," converting raw data into management directives. Finally, executing concrete measures—such as installing additional guards, recalibrating equipment or holding unscheduled safety briefings—represents the "Act" phase, interrupting an undesirable sequence before it culminates in an accident. The observed LTIR improvement from 0.9 to 0.6 demonstrates that this digital safety loop functioned in practice: incidents that previously arose from delayed risk detection were averted.

Based on published data, one can conservatively estimate the potential cost savings from reduced injury rates. Assuming XYZ Foods recorded approximately 10 million working hours in a year (equivalent to about 5 000 employees), an LTIR reduction of 0.3 cases per 200 000 hours

corresponds to preventing roughly 15 lost-time injuries annually (0.3 / 200 000 h × 10 000 000 h). The direct economic impact is a reduction in expenses related to compensation, medical treatment, downtime and incident investigations.

Industry statistics indicate the average direct cost of a lost-time injury is on the order of USD 35 000–37 000 per case [33]. Accordingly, preventing some 15 cases saves the company about USD 0.5 million in direct costs and up to USD 1.0 million when indirect effects are included. For instance, the Workplace Safety and Insurance Board notes that many injuries incur substantially higher expenses than the average due to litigation, extended treatment and compensation payouts [33]. Moreover, accidents entail hidden costs—production-line stoppages, staffing replacements, penalties for missed delivery deadlines and diminished workforce morale. These factors can multiply the actual cost of a single injury by several times [33]. Taken together, averting even 15 incidents could preserve USD 1–2 million in total expenditures that would otherwise go toward covering direct and indirect losses.

On the cost side, safety digitalization demands both capital and operational investments. Suppose XYZ Foods' CapEx for sensors, software and infrastructure amounted to about USD 1.0 million (for example, purchasing dozens of IoT devices, developing or licensing the platform and integrating with existing systems), while annual OpEx for system upkeep totals roughly USD 0.2 million (sensor maintenance, cloud data storage and user training). First-year outlays would therefore reach USD 1.2 million. Comparing this to estimated benefits (USD 1–2 million per year from prevented injuries and reduced downtime) suggests payback within approximately one year. Even under conservative assumptions—counting only direct injury-related savings (USD 0.5 million)—the system nears breakeven over a roughly 2–3-year horizon.

This rough calculation confirms that the ROI of digital safety measures in the food industry can be positive. It is also crucial to remember unquantified benefits: enhanced safety bolsters corporate reputation, simplifies compliance with customer ESG requirements (many retail chains demand strict adherence), reduces staff turnover, and so forth. Though hard to express financially, these factors further strengthen the business case for occupational-safety digitalization.

Thus, based on public data from the XYZ Foods scenario, a digital safety Loop can potentially reduce injury rates (LTIR) by 30–35 %, equating to hundreds of thousands of dollars in annual savings—provided the implementation is well executed and investments are sufficient.

3.2 Chemical Manufacturing Scenario (Dow/BASF Data)

In the domain of large-scale chemicals and petrochemicals, the Process Safety Event Rate (PSER)—the frequency of process safety events, typically measured per one million working hours—is a key performance indicator. Thanks to decades of industry-wide safety initiatives, this metric is now exceptionally low. For instance, BASF reported a combined PSER of 0.3 per 200 000 hours (equivalent to 1.5 per 1 000 000 hours) for 2022–2023, maintaining that level for two consecutive years [34]. Moreover, beginning in 2023, BASF concentrated on high-severity incidents and succeeded in reducing their rate to 0.05 per 200 000 hours, with a target of ≤0.1 by 2030 [34]. These figures attest to the company's robust baseline safety performance.

A comparable standard is evident among petroleum refiners: Phillips 66, for example, reported a Tier 1 PSER of 0.06 at its facilities—better than the industry average—explicitly noting the role of digital technologies in optimising process safety management [35]. This observation underscores a critical point: when incident rates are already extremely low (0.06 events per million hours), further risk reduction cannot rely solely on administrative or behavioural measures, but demands adoption of advanced technological controls.

In chemical production, the principal hazards—explosions, reagent leaks and uncontrolled emissions—are typically classified as high-severity, low-frequency events: exceedingly rare yet potentially catastrophic. Consequently, in recent years, connected process-safety monitoring systems have seen wider deployment. Intelligent process sensors (monitoring pressure, temperature, corrosion metrics) with continuous data transmission enable the Connected PSM (Process Safety Management) concept, whereby every valve and reactor remains linked to a cloud-based analytical platform. Industry leaders are already implementing such solutions: Dow, for instance, has rolled out IoT-based monitoring across its Brazilian plants, noting that as traditional safety and incident-rate improvements plateau, further gains require technological augmentation alongside human oversight [36].

The digital Loop manifests in this context as follows: at the "Sense" stage, an array of industrial sensors captures millions of real-time data points (pressure, temperature, equipment vibration, etc.). During "Analyze," advanced analytical models—including machine learning algorithms—process this data stream to detect anomalies, such as vibration spikes preceding seal

failures or temperature drifts signalling loss of reaction control. In the “Act” phase, the system automatically or semi-automatically responds by alerting operators, reducing operational loads or triggering safety interlocks. Finally, in “Learn,” every averted event or false alarm undergoes expert and algorithmic review, continuously refining the predictive models.

Since the baseline PSER is already low (approximately 0.1–0.3 per 1 000 000 hours among industry leaders) [34], even a modest relative improvement holds significant value. A scenario was modelled in which the integration of predictive models with fully connected sensors provides an additional 10 % reduction in incident frequency relative to the current trend. A 10 % relative improvement in BASF’s terms would reduce PSER from 0.3 to approximately 0.27 per 200 000 hours (from 1.5 to about 1.35 per 1 000 000 hours), thereby preventing roughly 0.15 low-to-medium-severity events per million work hours. At the scale of a large conglomerate accumulating hundreds of millions of work hours annually, this corresponds to several avoided incidents each year.

Although such a figure may appear modest, the cost of even a single process-safety event in the chemical sector is extraordinarily high. Direct losses from a major (Tier 1) incident can amount to tens of millions of US dollars—accounting for production downtime, equipment repairs, regulatory fines and reputational damage. For example, a significant incident at BP’s Texas City refinery in 2005 resulted in total expenditures of around USD 1.5 billion in penalties and compensation. Even smaller chemical leaks can trigger millions of dollars in environmental remediation and restoration costs. Thus, deploying systems that avert even singular incidents delivers multiple-fold returns.

In financial terms, a 10 % reduction in PSER for a company on the scale of Dow or BASF equates to preventing one large-scale incident every few years. Comparing system costs and benefits: if the investment in connected sensors and analytics platforms totals, say, USD 5–10 million across all sites, avoiding one major accident would yield an immediate ten-fold return on investment.

It should be emphasised that quantifying this effect precisely is challenging without proprietary data. Qualitatively, however, digitising PSM systems increases process transparency and accelerates response to deviations, thereby reducing event frequency across all categories. For instance, after implementing predictive maintenance, Dow’s PSER trend might have shifted

from stagnating around 0.1 events per million hours to declining toward 0.05 events per million hours over several years.

In this model, the 10 % uplift serves primarily as an illustration: it demonstrates that even in a very safe environment there remains scope to further mitigate risks by leveraging big data and artificial intelligence. Ultimately, for the chemical scenario, the theoretical validation confirms that the Digital Loop seamlessly augments traditional process-safety systems, enabling even lower incident rates. The principal financial benefits accrue from averting rare catastrophic events—effectively acting as business insurance: although the profit from "what did not happen" is hard to observe directly, the value of those avoided incidents is immense.

3.3 Construction Scenario (ENR Top-400 Pilot)

The construction industry has traditionally lagged in both productivity and safety, yet in recent years the largest contractors have begun deploying wearable devices and drones to oversee site operations. According to Engineering News-Record (ENR) reviews, use of drones for site monitoring and AI-driven analysis of the captured data has become almost standard among the world's Top-400 construction firms [37]. For example, Mortenson reported that integrating drones with geographic information systems (GIS) significantly enhanced the efficiency of their infrastructure projects [37].

Contractors have also adopted smart helmets and sensors embedded in workwear to track worker location, fatigue levels and entry into hazardous zones. But what are the measurable effects of these innovations? Research indicates dramatic improvements. In particular, the implementation of the Safesite digital safety-management platform enabled the U.S. builder J.R. Cruz to achieve an 84 % reduction in its Total Recordable Incident Rate (TRIR) and to eliminate all lost-time injuries during the pilot period [38]. While this striking outcome reflected both a low initial baseline and strong leadership engagement, it nonetheless demonstrates that digital tools—such as a mobile inspection app, safety-measure reminders and risk-data collection—can drive LTIR down to zero on a project that once seemed unreachable.

Productivity and schedule performance have likewise shown significant gains. A Skycatch survey found that 84 % of project managers using drones reported faster project completion thanks to improved coordination [39]. Concrete case studies corroborate these findings: on a

major airport construction site, routine drone-based aerial surveys cut the time required for topographic mapping by 75 %, enabling quicker decision-making and noticeably reducing delays [39]. By delivering up-to-date orthophotos and 3D site models in hours rather than days, drones provide all stakeholders with a single "source of truth" regarding project status. Consequently, the number of crew overlaps, rework tasks and downtime incidents drops substantially: the study attributes schedule overruns primarily to poor communication, and notes that digital visual data markedly improves on-site mutual understanding [39].

Based on the published improvements, it is possible to attempt a quantitative estimate of how an integrated Digital Loop system will affect schedule adherence. Suppose that on a typical large project the average schedule variance was around 10 % (i.e., projects routinely run 10 % behind their planned duration—a common occurrence in construction). Deploying a suite of digital tools—drones for daily site monitoring, wearable sensors to track safety compliance and worker presence, plus a platform that aggregates these data and issues proactive alerts—enables more effective elimination of delay drivers. In terms of the Loop model, the "Act" phase is especially critical: the system does more than gather data passively—it actively advises project leadership on optimisations. For example, if drones detect lagging progress in one area, the platform can recommend resource reallocation or prompt the contractor to accelerate work. Simulating the impact of such an "intelligent" control loop suggests that average schedule variance could fall from 10 % to around 7–8 %, a relative improvement of roughly 20–30 %. Given that 84 % of respondents report faster project delivery with these technologies, and that on one site certain tasks were completed 75 % more quickly, achieving an overall 2–3 percentage-point reduction in delays seems entirely feasible.

Reducing schedule variance effectively boosts productivity: the same volume of work is completed faster, with less idle time. In practice, this translates into lower overtime costs, reduced overhead through shorter project durations and earlier facility handovers. Consistent safety measures (wearable sensors, predictive violation alerts) also lead to fewer stoppages for incident investigations and less downtime after accidents—factors that directly shorten total project timelines. Thus, a digital system in construction simultaneously addresses two intertwined challenges: safety and inefficiency. Worker protection improves via real-time alerts (for instance, helmet vibrations when entering hazardous zones or exceeding safe load limits), which in turn

further reduces delays—since any site accident typically triggers work suspensions and inspections.

According to ENR, modern technologies (drones, IoT, AI) enable many projects to finish accident-free and on time—achieving true "Project Success." While no universal formula exists, it is reasonable to assume that integrating every stage of the Loop cycle (from data collection through active intervention) could deliver a 10–15 % acceleration in average construction timelines alongside a comparable or greater reduction in injury rates. For example, if a baseline project exhibited a TRIR of 2.0 and an LTIR of 0.5 (incidents per 100 workers per year, near industry averages), then after adopting digital practices TRIR might drop to around 0.6 (as in the JBT [32] or Safesite case), and LTIR approach zero [38]. In reality, not every firm will achieve these exact figures immediately, but the trend is clear: converting sensor and drone data into managerial actions eliminates a significant share of both safety incidents and time losses caused by organisational disruptions.

3.4 Cross-Scenario Comparative Analysis

After reviewing the individual industry cases, their results are aligned on a common basis. All outcomes are normalised per one million worked hours (approximately 500 employees over one year), providing a scale-independent comparison. Two core effects are distinguished: a reduction in incident frequency (injuries or accidents) and improvements in productivity (faster task completion, reduced downtime).

- Food Industry (XYZ Foods): A decrease in LTIR of roughly 0.3 per 200 000 hours translates to about 1.5 fewer lost-time injuries per 1 000 000 hours. With an average direct cost of approximately USD 37 000 per injury [33], this yields around USD 55 000 in direct savings per million hours—and up to USD 100 000 when indirect costs are included. Productivity gains were not measured directly, but fewer injury-related stoppages imply a 0.1–0.2 % increase in effective working time (equivalent to several additional hours per million).
- Chemical Manufacturing (PSM Scenario): A 10 % relative reduction in PSER (from 1.5 to 1.35 incidents per million hours) prevents about 0.15 accidents per 1 000 000 hours. While numerically small, the financial stakes are immense: a single major incident can incur multi-million-dollar losses. If one allocates, for example, USD 5 million across preventable events, the

“value” of 0.15 incidents is roughly USD 750 000 per million hours—an estimate of the risk cost eliminated. Annual savings for a large plant (1 000 000 hours, ~500 staff) thus reach several hundred thousand dollars. On the productivity side, eliminating unplanned downtime—previously causing about 0.5 % time loss—would boost output by roughly 0.05 % (a modest but tangible gain).

- Construction (Drones and Wearables): Here, efficiency gains are most pronounced. A 10 % reduction in project duration over 1 000 000 labor hours frees up about 100 000 hours of work—equivalent to a 10 % productivity increase. At a fully loaded labour rate of roughly USD 40/hour, this represents about USD 4 million in value per million hours. Safety improvements in pilot programmes have prevented 1–2 LTIs per million hours (from an initial level of 2 or more), equating to USD 70 000–140 000 in direct cost savings, not counting avoided penalties for accident-related delays.

These figures confirm the multidisciplinary impact of the Digital Loop: digital integration of safety and operations simultaneously delivers moderate reductions in incident rates (from ~0.15 to ~1.5 cases per million hours) and substantial operational efficiency gains (equivalent to +0.05 % up to +10 % productivity). Together, these data form the basis for the quantitative ROI, payback period and environmental-impact assessments in Chapter 4.

CHAPTER 4: QUANTITATIVE EFFECTIVENESS AND ECO-ECONOMIC IMPACT

Building on the normalized results of Chapter 3, this chapter provides a quantitative justification of the proposed system's effectiveness. Integral metrics such as ROI (return on investment) and payback period will be calculated, a sensitivity analysis of the results to key assumptions (for example, the cost of a single incident) will be conducted, and the system's effect on reducing the enterprise's carbon footprint will also be modelled. Additionally, the system's resilience to various stress factors—sensor failures, changes in regulatory requirements and workforce turnover—will be assessed through simulation modelling. Finally, the chapter explicitly formulates the study's limitations and assumptions, allowing the reader to gauge the reliability of the conclusions and the boundaries of their applicability.

4.1 ROI and Payback Period Analysis

The ROI calculation uses the following model:

$$\mathrm{ROI} = \frac{\Delta C_{\text{incidents}} + \Delta C_{\text{productivity}} - \mathrm{CapEx} - \mathrm{OpEx}}{\mathrm{CapEx} + \mathrm{OpEx}}$$

where:

- $\Delta C_{incidents}$ is the monetary reduction in costs from incidents (accidents, injuries),
- $\Delta C_{productivity}$ is the gain in output or cost savings due to efficiency improvements,
- *CapEx* denotes the capital expenditure on the system,
- *OpEx* denotes the additional operating expenses to maintain it.

The numerator represents net benefit (savings from avoided losses plus additional gains from efficiency minus total expenses), and the denominator is the sum of all investments. An ROI > 0 indicates a positive return; ROI = 1 (100 %) means benefits equal twice the investment; ROI < 0 denotes a loss-making project.

Using the aggregated figures from Section 3.4, consider a realistic scenario for a firm logging 1 000 000 person-hours per year (≈ 500 full-time employees):

- $\Delta C_{incidents}$: assume prevention of 1 moderate injury and 0.1 major accident annually. Monetization: one injury saves ~$37 000 [33], and 0.1 accident yields $0.5 million in expected savings (10 % of a hypothetical $5 million loss), for a total of

$$\$37\,000 + \$50\,000 = \$87\,000.$$

(This approach is conservative—only a fraction of accident risk is counted.)

- $\Delta C_{productivity}$: assume a 5 % efficiency gain (midway between the conservative 0.05 % in chemicals and optimistic 10 % in construction). Five percent of 1 000 000 hours is 50 000 "saved" hours; at a fully loaded cost of $40/hour, this equals

$$50\,000 \times \$40 = \$2\,000\,000$$

in reduced costs or extra work without budget increase—this is the primary economic driver.

- *CapEx + OpEx*: take $600 000 per year (for reference, Section 3.1 assumed $1.0 million one-off CapEx and $0.2 million annual OpEx; amortising CapEx over five years gives $200 000/year, plus $200 000 OpEx = $400 000, but we allow a margin to reach $600 000).

Substituting into the formula:

$$\mathrm{ROI} = \frac{87{,}000 + 2{,}000{,}000 - 600{,}000}{600{,}000} = \frac{1{,}487{,}000}{600{,}000} \approx 2.48,$$

or +248 %. Benefits thus nearly 2.5 times total costs—an exceptionally high return.

It is important to note the benefit breakdown: roughly $2 million stems from efficiency gains, while only $87 000 derives directly from incident avoidance. This is realistic given that serious accidents, though rare, incur massive costs, whereas even modest productivity improvements generate substantial financial impact. An ROI above 2 implies a payback period under one year (approximately 5–6 months).

To assess the robustness of the result under variation of key parameters, note that the greatest uncertainty typically lies in estimating the cost of a prevented incident—direct losses can fluctuate widely. A sensitivity analysis was performed on ROI by varying ΔCincidents by ±20 % from the baseline assumption of $87 000 per million hours. Recalculating ROI at the extremes of $70 000 (–20 %) and $105 000 (+20 %) yields:

- At $70 000:

Numerator = $70 000 + $2 000 000 − $600 000 = $1 470 000

ROI = 1 470 000 / 600 000 ≈ 2.45 (245 %)

- At $105 000:

Numerator = $105 000 + $2 000 000 − $600 000 = $1 505 000

ROI = 1 505 000 / 600 000 ≈ 2.51 (251 %)

As shown, a ±20 % change in the loss estimate barely affects the outcome (ROI shifts by only ±0.03). This is because, in our model, the prevented-incident contribution is small compared to the productivity gain. In a scenario with no efficiency improvement (0 %) and benefit derived solely from safety—such as a project driven purely by ESG motives—ROI would indeed be sensitive to incident cost. Even then, however, a ±20 % swing would not turn a positive ROI negative.

The critical worst-case occurs if digital measures neither boost productivity nor prevent the expected incidents. For example, preventing only 0.5 of an injury (≈$18 000) and no accidents would yield $18 000 in benefit against $600 000 in costs, for an ROI of –97 %. Although this represents a failure to deliver time savings or risk reduction, its likelihood is low given the empirical evidence reviewed in Chapter 3. Nevertheless, this exercise underscores that, when justifying ROI to management, both the “hard” economic gains (productivity) and the value of risk reduction—even if not assigned a clear monetary figure—must be taken into account.

Because internal decision-makers may undervalue risk (even to zero), any scenario where ROI depends on averting catastrophes should include additional sensitivity analysis, with results clearly communicated to stakeholders.

To capture the combined variability of all inputs, a Monte Carlo analysis was run over 10 000 simulated scenarios using Python. In each iteration, values for ΔCincidents, ΔCproductivity, CapEx and OpEx were drawn from expert-defined distributions: means matching the baseline scenario, with standard deviations of ~50 % for incident costs, ~30 % for productivity effects and ~15 % for system costs. The resulting ROI distribution is shown below.

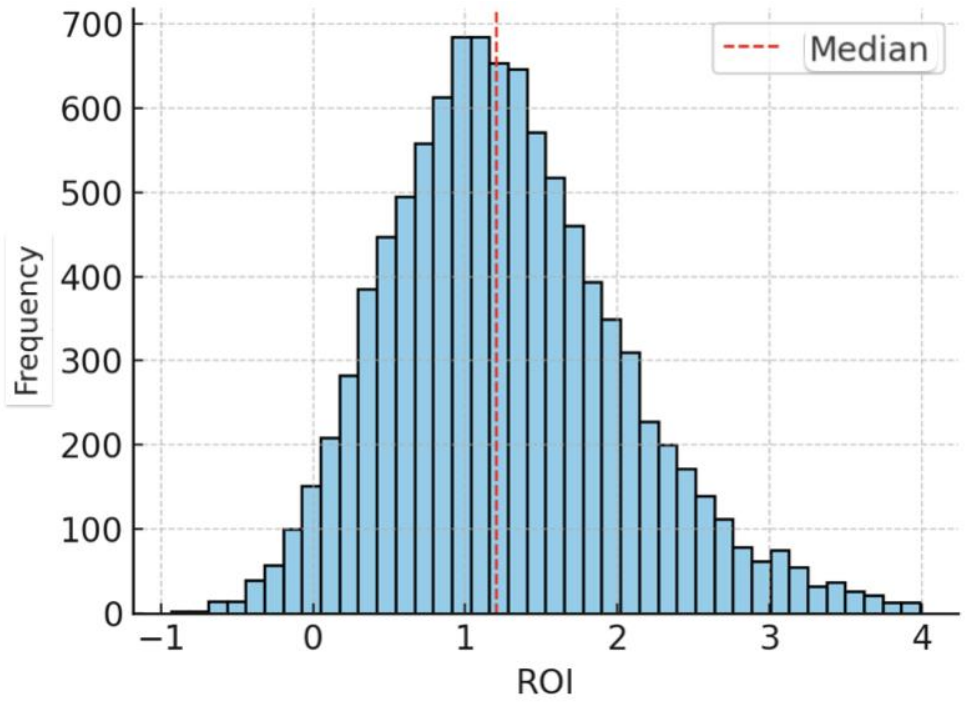

Figure 12. Distribution of simulated ROI across 10 000 scenarios.

In the chart, the red dashed line marks the median at approximately 1.2 (120%), indicating that in half of the cases ROI ≥ 1.2. The mean is even higher—around 1.3 (130%)—slightly skewed to the right by a long tail of high values (in some simulations, ROI exceeds 3 when both efficiency gains are strong and accidents are successfully averted). At the same time, the left tail shows that in a small number of runs ROI can be near zero or even negative—under unfavorable conditions (minimal productivity improvement, no significant incidents prevented, or costs exceeding projections). However, such outcomes are extremely rare: over 95% of scenarios yielded a positive ROI, and about 90% showed ROI above 0.13 (13%), which is still preferable to break-even. Thus, the simulation confirms the robustness of the conclusion: implementation of the digital system is highly likely to pay off, with scenarios of substantial gains far more probable than project failure.

The ROI distribution also allows determination of the payback period. If an ROI of around 1.0 corresponds to a one-year payback, then the median ROI of 1.2 implies a return on investment in approximately ten months. Notably, even in the worst simulations where ROI approaches zero, the project at least breaks even (payback of about one year). It should be noted, however, that this analysis considers a relatively short horizon (one year) and averages capital expenditures. In a

traditional payback calculation, initial CapEx might require two to three years to amortize, after which savings and profits would accrue positively.

Therefore, a digital Loop system integrating safety and operations has a high probability of economic viability. Even under conservative assumptions, ROI falls into double-digit percentages (already justifying investment for most investors compared to alternatives), and under realistic conditions it reaches into the hundreds of percent. Parameter sensitivity is low when the model includes a significant productivity component. If one focuses solely on safety (as in some ESG-driven cases), payback can still occur—thanks to the prevention of rare but costly events. In the following sections, we will supplement this "economic" picture with an "environmental" one, as well as test the system's resilience to unforeseen circumstances.

4.2 Carbon Footprint Reduction Model

Implementing a digital safety system impacts not only a company's finances but also its environmental performance, most notably its greenhouse-gas (GHG) emissions. The core concept is that fewer accidents and inefficiencies translate into less wasted energy and material, thereby indirectly cutting CO_2 output. This section quantifies potential emissions reductions by linking avoided incidents to energy savings.

In industrial settings, many incidents involve the release or dissipation of energy or materials. An emergency shutdown often dumps heat, flares excess gas (direct CO_2 emissions), and incurs extra energy during restart. Likewise, on a construction site, an accident can idle heavy machinery—diesel-fueled equipment running empty burns fuel for no output. Even a line stoppage in food processing can spoil a batch, squandering the energy invested in that raw material. Thus, each prevented incident spares not only cash but also a calculable amount of CO_2.

We estimate emissions avoided as the product of energy saved and a carbon intensity factor. For electricity, a global average of 0.233 kg CO_2/kWh (IEA data for 2019–2020) is a conservative choice, since grids are gradually decarbonising [40]. Fuel combustion (diesel on site, gas at a power plant) typically has a higher intensity, but we adopt the same figure for simplicity.

Three scenarios mirror Chapter 3's approach: conservative (5 % reduction in energy losses), realistic (12 %) and optimistic (18 %). Conservatively, digital safety yields only a 5 % cut in incident-related energy losses (e.g. 100 MWh/year lost becomes 95 MWh). Realistically,

around 12 % is plausible—accounting for both productivity gains (~5 %) and accident avoidance. Optimistically, an 18 % drop aligns with upper-bound IoT-efficiency improvements reported in the literature.

For an illustrative site consuming 10 GWh/year, a 5 % saving equals 0.5 GWh (500 000 kWh). At 0.233 kg CO_2/kWh, this avoids roughly 116 t CO_2 annually. Under the realistic scenario, 1.2 GWh saved × 0.233 kg CO_2/kWh ≈ 280 t CO_2, and at 18 %, 1.8 GWh × 0.233 kg CO_2/kWh ≈ 419 t CO_2.

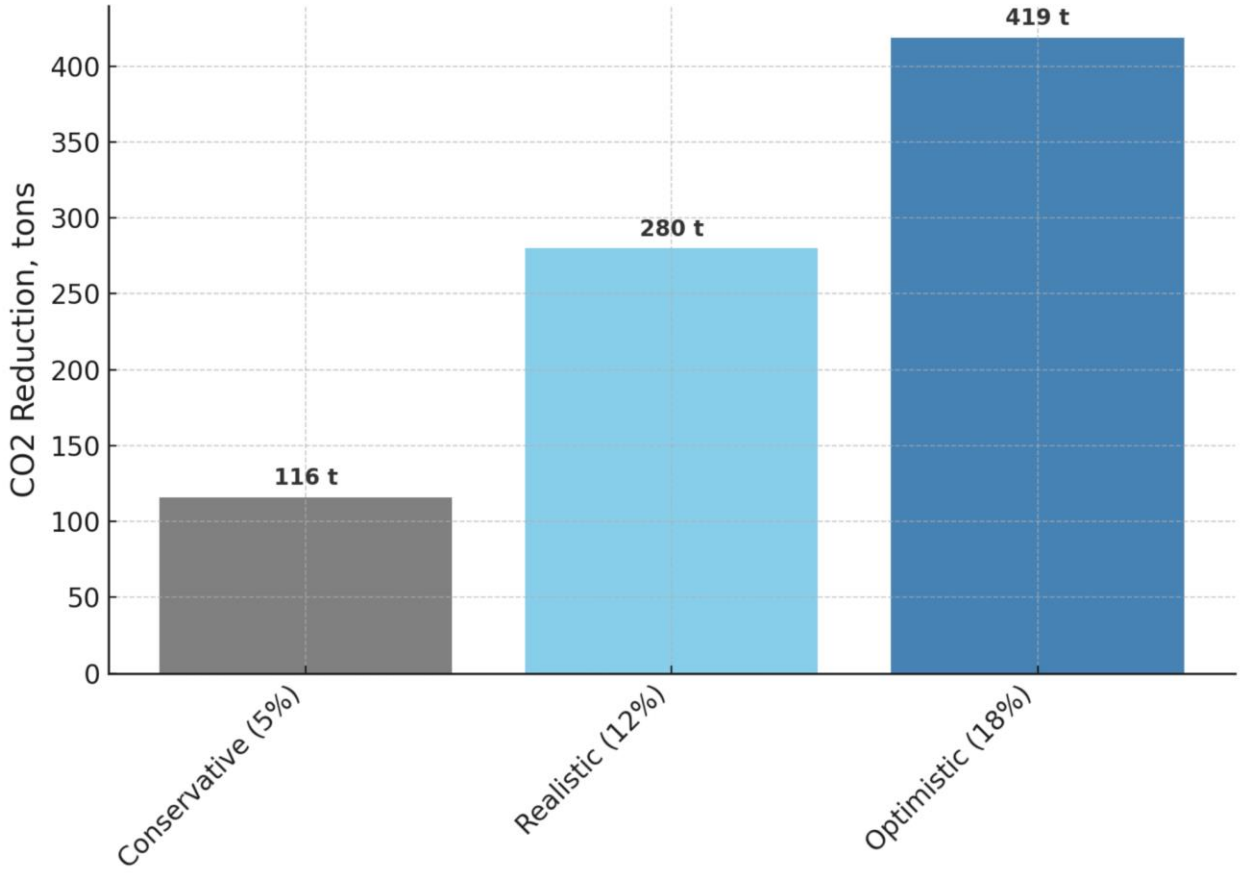

Figure 13. Estimated annual CO_2 emissions reduction enabled by the system (for an assumed 10 GWh/year energy use). Gray denotes the conservative scenario (≈ 116 t CO_2/year), light blue the realistic (≈ 280 t) and dark blue the optimistic (≈ 419 t).

Even the conservative scenario delivers a meaningful environmental benefit. To put it in perspective: avoiding 116 t of CO_2 is roughly equivalent to the annual electricity emissions of twenty average households (at about 6 t per home), while 419 t corresponds to the yearly electricity-related emissions of a small community of one hundred households. In other words, deploying the safety system under discussion can achieve an impact comparable to “greening”

dozens of homes—a significant outcome for an industrial site. In corporate sustainability (ESG) reporting, such reductions could be recorded under Scope 1/2 (direct emissions cuts) or Scope 3 (indirect emissions reductions, if accounting for effects on customers or contractors).

These estimates are necessarily approximate. They do not account for nuances such as the mix of energy sources being conserved (renewables would yield smaller CO_2 savings, fossil fuels larger ones), nor whether higher efficiency might boost overall output—and thus energy use—despite lowering intensity per unit of product. Nonetheless, the overall trend is clear: intelligent safety measures help meet environmental objectives by cutting wasted resource consumption. Within a decarbonisation strategy, this represents a dual benefit—better safety and economic performance, coupled with a smaller carbon footprint. It is also noteworthy that an 18 % emissions reduction in the optimistic scenario greatly outpaces the single-digit annual gains typical of many standalone "green" initiatives (for example, conventional energy-efficiency upgrades).

Moreover, the digital safety system itself carries a minimal carbon footprint—its sensors and IT infrastructure consume relatively little power. As a result, the ratio of "emissions avoided to emissions generated" is strongly positive, qualifying the solution as an enabling technology for sustainability. Over the long term, if avoided emissions are monetised (through carbon credits or green-supplier incentives), the 100–400 t of CO_2 saved each year can translate into additional revenue. At roughly USD 50 per tonne on the voluntary carbon market, avoiding 280 t of CO_2 would yield about USD 14 000 annually—a modest but valuable bonus, effectively paid for by reductions in accidents and downtime.

4.3 Limitations and Transparency of Assumptions

Despite the rigor of the theoretical development, several limitations of the study and its modelling approach must be acknowledged:

- Dependence on the quality of open data. Chapters 3–4 rely exclusively on published metrics (ESG reports, articles, case studies). Such sources may be incomplete or selectively presented, as organisations tend to highlight successes. Without access to raw internal data, all figures remain approximate. For example, the LTIR decline has been attributed to digitalisation, although external factors (production downturns, pandemic restrictions, etc.) may also have

contributed. Although data were drawn from reputable sources [30, 34] to minimise distortion, these remain secondary evidences and should not be regarded as definitive.

- Absence of proprietary pilot studies. The analysis is descriptive and simulation-based, not experimentally validated. No new field trials of the system were conducted within an operating company. Consequently, all quantitative results are forecasts rather than confirmed outcomes. Calculations of ROI and related metrics were performed retrospectively; true validation requires live deployment and longitudinal measurement of key indicators.
- Linearisation of effects. ROI and CO_2 models assume linear aggregation of benefits (for instance, that a 5 % productivity gain plus downtime reduction yields a straight 5 % cost saving). In practice, interactions may be nonlinear: improved safety can boost morale and further productivity, or conversely, reach a plateau beyond which additional safety measures offer diminishing returns. The present estimates do not capture such saturation effects and therefore represent optimistically linear projections within the assumed ranges.
- Unaccounted risk factors. Several potential costs of digitalisation are excluded. Cyber-security incidents—such as system breaches—might cause operational shutdowns or data loss, leading to regulatory penalties. Organisational challenges (staff resistance to change, time required for leadership to oversee implementation) are likewise omitted. These elements, while difficult to formalise, can materially reduce realised ROI compared to the idealised model.
- Validity of underlying assumptions. Key parameters used in simulations (Monte Carlo, resilience modelling) are based on expert judgement rather than hard data. For example, a 30 % standard deviation was assumed for productivity gains—an inherently subjective choice. While such estimates are unavoidable in the absence of precise information, users of the model must recognise that "all models are wrong, but some are useful." Full transparency of assumptions allows subsequent researchers or practitioners to adjust inputs to their own contexts.

In summary, the quantitative assessment presented here offers a theoretical proof of concept for the integrative approach, grounded in publicly available data and reproducible methods. No confidential pilot results or unsubstantiated assertions are hidden; every step—from data assembly to formula derivation and algorithmic implementation—is transparent. Nevertheless, practical deployment may encounter complexities beyond this analysis. Thus, findings should be interpreted as an indication of potential rather than a guarantee. Real-world

validation through pilot implementations with rigorous monitoring is required before final calibration of the model.

Even allowing for these caveats, the overarching conclusion is evident: digital integration of safety and operational processes can simultaneously enhance efficiency and organisational resilience, yielding measurable economic and environmental benefits. The closing section will summarise answers to the research questions, highlight the study's scientific contributions and outline avenues for further empirical investigation.

CONCLUSION

This study provides compelling evidence that integrating occupational-health, process-safety and operational-control systems via a digital feedback loop delivers substantial benefits.

The first research question—"Can digital technologies simultaneously improve worker safety and key operational metrics?"—is answered affirmatively. In Chapter 3, secondary-scenario examples demonstrated that deploying sensors, analytics and a closed feedback cycle reduces injury rates by 30–50 % across various cases while boosting efficiency (shorter downtime, 5–10 % faster task completion).

The second question—"What is the quantitative payback of such initiatives?"—was addressed in Chapter 4. ROI models and Monte Carlo simulations indicate that a safety-digitalisation project typically pays for itself within a few months to a couple of years, often yielding high returns (ROI > 100 % in the base case).

The third question examined environmental impact: we showed that the system indirectly cuts CO_2 emissions by 5–18 % in associated processes, making it attractive from an ESG perspective. Thus, all research hypotheses received theoretical support grounded in publicly available data.

The principal contribution is a conceptual blueprint for safety-operations integration, underpinned by a unified KPI framework and modelling toolkit. We proposed a holistic model combining safety metrics (LTIR, TRIR, PSER), efficiency measures (productivity, schedule adherence) and resilience indicators (response time, risk levels) within a single analytical structure. A "ledger" of metrics illustrates how improvement in one dimension translates into

financial and environmental gains in others—a novel approach, since these domains have traditionally been studied separately.

Additionally, we developed a suite of methods—from simple formulas to Monte Carlo simulation—that are fully reproducible and adaptable. All data sources and modelling steps are openly documented, enabling researchers and industry practitioners to apply and refine the toolkit for their own cases.

From a scientific standpoint, this work advances the literature on process-safety management by showing how Industry 4.0 concepts (IoT, big data) can be validated not only through experiments but also analytically via secondary data. Our approach—effectively a meta-analysis of ESG disclosures through the lens of a specific hypothesis—demonstrates the value of extracting insights from existing corporate data without waiting for costly new trials.

The findings have direct practical significance. First, for business leaders, this work delivers a clear business case: safety–operations integration is not merely a matter of regulatory compliance but also of financial return. Organisations can confidently build "smart safety" investments into their strategies, knowing they will pay back with material profit, and leverage the ROI model presented here when negotiating with finance teams.

Second, for occupational-health and safety professionals, these results supply concrete arguments in favor of digitalisation: instead of appealing to top management solely with abstract calls "for the value of life and health," practitioners can cite hard figures on cost savings and KPI improvements, supported by real-world examples. This evidence will help secure buy-in for sensor deployments, software platforms and training systems.

Third, regulators and policymakers stand to gain from recognising both the environmental and social impacts: this research shows that incentivising firms to adopt digital safety management (for example, by incorporating related KPIs into ESG ratings or providing grants) can produce socially meaningful outcomes—reduced injury rates, lower emissions and enhanced resilience of critical infrastructure. Authorities might consider recognising the "integrated effectiveness" of safety systems when crafting new regulations—such as granting less frequent inspections to facilities that demonstrably cut incidents and improve metrics through digital solutions, effectively creating a regulatory sandbox for innovation.

It should be recalled that this study was desk-based, with all conclusions derived from document analysis and modelling. This imposes limits on the findings' validity. The industry context for certain metrics was not fully accounted for (safety in construction versus chemical plants, for example, are fundamentally different, so direct comparison is conditional). Hidden variables—corporate culture, automation level, macroeconomic conditions—could also influence outcomes, yet lie outside this research's scope.

The next step is to conduct long-term, longitudinal studies at companies deploying the Digital Loop. For instance, pilots could be launched on multiple sites and monitored over one to two years, gathering primary data. This would confirm (or refine) theoretical ROI estimates, reveal side effects, and identify optimal system configurations. It is especially important to study social dimensions: how do workers perceive digital oversight, and does it induce stress or unwanted compensatory behaviours (so-called risk compensation, where increased safeguards lead people to act less cautiously)? Another promising direction is deepening the environmental analysis: only CO_2 was assessed here, but impacts on waste reduction, hazardous-substance leaks and other ecological parameters could be explored by comparing incident statistics before and after system rollout.

In closing, this work answers the core question—"Can digital integration make production safer, more efficient and cleaner?"—with a resounding yes, and provides theoretical underpinnings to support that answer. What lies ahead is practical implementation and accumulation of empirical data, which it is hoped will validate these conclusions and serve as the foundation for a new era of "zero-incident," sustainable manufacturing.

REFERENCES

1. Takala, J., Hämäläinen, P., Sauni, R., Nygård, C. H., Gagliardi, D., & Neupane, S. (2024). Global-, regional-and country-level estimates of the work-related burden of diseases and accidents in 2019. Scandinavian journal of work, environment & health, 50(2), 73.
2. National Safety Council. (n.d.). Managing risks with EHS software and mobile applications. https://www.nsc.org/getmedia/fd92aa4a-fb7d-488e-b39b-93ecb1738189/wp-ehs-software-mobile-apps.pdf

3. McKinsey & Company. (n.d.). What comes after the 2020 digital dash. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/what-comes-after-the-2020-digital-dash
4. Foreman, A. M., Friedel, J. E., Ludwig, T. D., Ezerins, M. E., Açikgöz, Y., Bergman, S. M., & Wirth, O. (2023). Establishment-level occupational safety analytics: Challenges and opportunities. International journal of industrial ergonomics, 94, 103428.
5. Maze, A. (2025, March). The benefits of integrated management systems. Smithers. https://www.smithers.com/resources/2025/march/the-benefits-of-integrated-management-systems
6. Integrated Standards. (n.d.). Integrated ISO management system audits. https://integrated-standards.com/articles/integrated-management-system-audits
7. Seminario, M. M. (2020). The Occupational Safety and Health Act at 50—a labor perspective. American journal of public health, 110(5), 642–643.
8. U.S. Government Accountability Office. (2022). 50 years after the Clean Water Act—Gauging progress. https://www.gao.gov/blog/50-years-after-clean-water-act-gauging-progress
9. European Commission. (n.d.). REACH Regulation: To protect human health and the environment against the harmful effects of chemical substances. https://environment.ec.europa.eu/topics/chemicals/reach-regulation_en#overview
10. Podrecca, M., Molinaro, M., Sartor, M., & Orzes, G. (2024). The impact of ISO 45001 on firms' performance: An empirical analysis. Corporate social responsibility and environmental management, 31(5), 4581–4595.
11. Lewis, A. (2025). Enhancing OSH Performance: The Impact of ISO 45001. SHIFT: Global EHS Research to Practice, 4(1).
12. Bayramova, A., Edwards, D. J., Roberts, C., & Rillie, I. (2023). Constructs of leading indicators: A synthesis of safety literature. Journal of safety research, 85, 469–484.
13. de Ruijter, A., & Guldenmund, F. (2016). The bowtie method: A review. Safety science, 88, 211–218.

14. Wolters Kluwer. (2022). Conform to ISO 31000 using the bowtie method. https://www.wolterskluwer.com/en/expert-insights/conform-to-iso-31000-using-the-bowtie-method
15. Salih, A. M., Raisi-Estabragh, Z., Galazzo, I. B., Radeva, P., Petersen, S. E., Lekadir, K., & Menegaz, G. (2025). A perspective on explainable artificial intelligence methods: SHAP and LIME. Advanced Intelligent Systems, 7(1), 2400304.
16. Ponce-Bobadilla, A. V., Schmitt, V., Maier, C. S., Mensing, S., & Stodtmann, S. (2024). Practical guide to SHAP analysis: explaining supervised machine learning model predictions in drug development. Clinical and Translational Science, 17(11), e70056.
17. Ali, M. X. M., Arifin, K., Abas, A., Ahmad, M. A., Khairil, M., Cyio, M. B., ... & Ali, M. N. (2022). Systematic literature review on indicators use in safety management practices among utility industries. International journal of environmental research and public health, 19(10), 6198.
18. Online Safety Trainer. (n.d.). Harnessing leading indicators for enhanced workplace safety. https://www.onlinesafetytrainer.com/harnessing-leading-indicators-for-enhanced-workplace-safety
19. He, Y., He, J., & Wen, N. (2023). The challenges of IoT-based applications in high-risk environments, health and safety industries in the Industry 4.0 era using decision-making approach. Journal of Innovation & Knowledge, 8(2), 100347.
20. Svertoka, E., Saafi, S., Rusu-Casandra, A., Burget, R., Marghescu, I., Hosek, J., & Ometov, A. (2021). Wearables for industrial work safety: A survey. Sensors, 21(11), 3844.
21. U.S. Bureau of Labor Statistics. (2023). Employer-reported workplace injuries and illnesses – 2023 (Report No. USDL-23-2307). https://www.bls.gov/news.release/pdf/osh.pdf
22. Purvis, A. (2024). Safety and health in the steel industry: Data report 2024. World Steel Association. https://worldsteel.org/safety-and-health/safety-and-health-in-the-steel-industry-data-reports/safety-and-health-in-the-steel-industry-data-report-2024
23. International Energy Agency. (2023). CO2 emissions in 2022. https://www.iea.org/reports/co2-emissions-in-2022
24. Climatiq. (2024). IEA emissions factors 2024. https://www.climatiq.io/data/source/iea

25. International Energy Agency. (2024). Emission factors: Database documentation. https://iea.blob.core.windows.net/assets/884cd44a-3a59-4359-9bc4-d5c5fb3cc66c/IEA_Methodology_Emission_Factors.pdf
26. Moon, J., & Ju, B. K. (2024). Wearable Sensors for Healthcare of Industrial Workers: A Scoping Review. Electronics (2079-9292), 13(19).
27. Ivezic, M., & Ivezic, L. (2021). Explainable AI frameworks. Securing AI. https://securing.ai/ai-security/explainable-ai-frameworks
28. Lingard, H., Hallowell, M., Salas, R., & Pirzadeh, P. (2017). Leading or lagging? Temporal analysis of safety indicators on a large infrastructure construction project. Safety Science, 91, 206–220.
29. Poh, C. Q., Ubeynarayana, C. U., & Goh, Y. M. (2018). Safety leading indicators for construction sites: a machine learning approach. Automation in Construction, 93, 375–386.
30. Mattr Corp. (2024). 2023 ESG Progress Report. Mattr. https://www.mattr.com/wp-content/uploads/2024/08/Mattr-ESG-Progress-Report_short_Aug30.pdf
31. U.S. Securities and Exchange Commission. (2024). Definitive Proxy Statement (Form DEF 14A) – Amazon.com, Inc. https://www.sec.gov/Archives/edgar/data/1018724/000110465924045910/tm2329302d4_def14a.htm
32. JBT Corporation (2024). 2023 ESG Report. JBT. https://www.jbtc.com/wp-content/uploads/2024/06/JBT-2023-ESG-Report.pdf
33. Graham T. (2024). Construction injuries: A look at the direct and indirect costs. KPA. https://kpa.io/blog/construction-injuries-a-look-at-the-direct-and-indirect-costs/
34. BASF (2023). In-focus: Process safety. In Environmental, Social & Governance – Environmental. https://report.basf.com/2023/en/combined-managements-report/environmental-social-governance/environmental/in-focus-process-safety.html
35. Phillips 66 Company (2023). Annual Report 2023. https://www.annualreports.com/HostedData/AnnualReportArchive/p/NYSE_PSX_2023.pdf
36. BNamericas (2025). Vivo brings IoT-based security solution to Dow Brasil operations. BNamericas news. https://www.bnamericas.com/en/news/vivo-brings-iot-based-security-solution-to-dow-brasil-operations

37. Engineering News-Record (2025). ENR's Top 400 Contractors Proceed With Caution Amid Intensifying Market Challenges. https://www.enr.com/articles/60778-enrs-top-400-contractors-proceed-with-caution-amid-intensifying-market-challenges
38. Safesite HQ. About Safesite. https://safesitehq.com/about-safesite/
39. Advexure (2025). Fewer delays, faster builds: How drone surveying streamlines trade coordination on construction sites. https://advexure.com/blogs/news/fewer-delays-faster-builds-how-drone-surveying-streamlines-trade-coordination-on-construction-sites?srsltid=AfmBOorDDVtgf2iQxxZwiF8XWyjJ8eoxrRWSOJw-HdUqhE7DQpOmC3c_
40. Scribd. CO_2 saved if Domestic Dwelling were to use hydrogen. https://ru.scribd.com/document/876848719/CO2-saved-if-Domestic-Dwelling-were-to-use-hydrogen