Unlocking Life-Saving Innovation: A Response to the Digital Omnibus Initiative from the Perspective of Occupational Safety and Health Professionals

ENSHPO's Contribution

Roberto Sammarchi, Attorney at Law, PhD, LLM
AIAS (IT) representative in ENSHPO

Executive Summary

The European Network of Safety and Health Professional Organisations (ENSHPO), a leading organisation representing occupational safety and health (OSH) professionals across Europe, has been actively engaged in the European Commission's Digital Omnibus initiative. This initiative is a crucial step towards modernising and adapting the regulatory framework to the rapidly evolving digital landscape. ENSHPO's primary advocacy focuses on the need for regulation that is both streamlined and highly effective. This dual focus is particularly crucial when considering the profound impact of Artificial Intelligence (AI) on workers' safety and health. Their main objective is to ensure robust protection of workers' physical and mental well-being within this transforming digital work environment. This report advocates for the reform of the EU's digital regulatory framework to remove barriers hindering the adoption of life-saving Occupational Safety and Health (OSH) technologies, proposing actionable changes to the AI Act, GDPR, and cybersecurity rules to foster innovation and achieve "Vision Zero" for workplace fatalities.

The emergence of AI and other digital technologies introduces a new spectrum of occupational risks that demand proactive attention. ENSHPO specifically highlights the need to address:

- Ergonomic Challenges: The introduction of new human-machine interfaces, AI-driven automation, and potentially altered working postures or repetitive tasks can lead to novel ergonomic issues, requiring updated risk assessments and preventative measures.
- Psychological Stress: Constant AI monitoring, algorithmic management, increased work intensity, blurring work-life boundaries, and potential job loss anxieties can

contribute to significant psychological stress among workers.

- AI Biases: Biases inherent in AI algorithms, if not carefully managed, can lead to discriminatory practices, unequal workload distribution, or even unsafe operational instructions, disproportionately affecting certain groups of workers.

To effectively address these challenges, ENSHPO advocates for an agile regulatory framework. Such a framework must possess the inherent flexibility to adapt rapidly to technological advancements while safeguarding fundamental worker rights and health. This approach necessitates a holistic consideration of the AI-driven work environment, encompassing:

- Physical Dimensions: Addressing traditional safety concerns alongside new risks related to collaborative robotics, remote-controlled machinery, and digitally augmented tasks. AI can be used to monitor workplaces in real-time, identify hazards, predict machine failures, and optimise processes to reduce exposure to physical risks.

- Psychological Dimensions: Recognising and mitigating the mental health impacts of AI, promoting psychological safety, and ensuring fair and transparent AI-driven management practices. AI can help identify patterns of stress or burnout, offer personalised support, and facilitate better communication to mitigate psychological stress.

- Social Dimensions: Ensuring that AI implementation fosters social inclusion, prevents new forms of discrimination, and upholds collective bargaining rights in the face of technological change. AI can be used to create more accessible training programmes, support worker reskilling, and enhance fairness in work distribution, contributing to a more inclusive and safer work environment.

The European Union stands at a critical juncture where technological advancement, particularly in Artificial Intelligence (AI) and the Internet of Things (IoT), offers an unprecedented opportunity to achieve the ambitious goal of "Vision Zero" for work-related fatalities and serious injuries. Technologies capable of predicting hazards, monitoring for unsafe conditions in real-time, and preventing catastrophic failures are no longer theoretical; they are a tangible reality. However, their widespread adoption is being significantly hampered by a complex, overlapping, and often ill-suited digital regulatory framework.

The OSH perspective

This report, submitted in response to the European Commission's "Call for Evidence" on the Digital Omnibus initiative, presents the collective perspective of Occupational Safety and Health (OSH) professionals. Our central analysis reveals that the EU's digital acquis—while rightly founded on principles of protecting fundamental rights, privacy, and security—inadvertently creates a significant "innovation chilling" effect on the very

technologies designed to save lives in the workplace. The cumulative burden of the AI Act, the General Data Protection Regulation (GDPR), the Data Act, and various cybersecurity reporting mandates creates a landscape of legal uncertainty, prohibitive costs, and practical impossibilities for both the developers and deployers of OSH technology.

The AI Act, with its broad, default classification of any worker-monitoring system as 'high-risk', fails to distinguish between invasive surveillance and protective monitoring. This blunt application imposes a disproportionate compliance burden, with high costs effectively barring smaller, specialized innovators from the market. Simultaneously, the GDPR's principle of data minimization is in direct conflict with the need for comprehensive datasets to train accurate and reliable safety-critical AI models. The conventional solution of data anonymization is often technically unfeasible or functionally crippling for real-time safety systems that must identify an individual in immediate danger.

This report puts forward a "Blueprint for Reform" with five core, actionable recommendations for the Digital Omnibus and the subsequent Digital Fitness Check:

1. Re-calibrate the AI Act's 'High-Risk' Definition for OSH: Issue clear guidance, in collaboration with EU-OSHA, to create a distinction between AI for surveillance and AI for protection, allowing life-saving systems to be assessed on their actual risk profile rather than by a catch-all default.
2. Establish OSH-Focused AI Regulatory Sandboxes: Leverage the AI Act's sandbox mechanism to create dedicated, supervised environments where innovators can test and validate OSH technologies in collaboration with regulators and social partners, accelerating market access and reducing legal uncertainty.
3. Mandate Harmonized Standards and Professional Certification: Commission OSH-specific harmonized standards for AI systems to provide a clear "presumption of conformity" with the law, drastically reducing compliance costs. Concurrently, establish a "Certified AI Safety Professional" (CAISP) certification to ensure genuinely qualified human oversight.
4. Clarify Data Governance Rules for Safety Systems: Provide explicit guidance that recognizes Privacy-Enhancing Technologies (PETs) as a primary means of GDPR compliance and confirms that an employer's legal OSH obligation constitutes a valid basis for processing necessary health data for safety purposes.
5. Unify Cybersecurity Reporting: Implement a single reporting portal for cybersecurity incidents with a risk-based, tiered timeline that prioritizes on-the-ground safety response over premature administrative reporting for non-critical vulnerabilities.

These proposals are not a call for deregulation but for smarter, more targeted, and technologically-aware regulation. By implementing these changes, the Commission can transform the digital acquis from a barrier into an enabler, fostering a European ecosystem

where the best technology is harnessed to make EU workplaces the safest in the world.

The Digital Frontier of Workplace Safety

The Untapped Potential

The landscape of occupational safety and health is on the verge of a profound transformation, driven by the convergence of Artificial Intelligence (AI), the Internet of Things (IoT), and wearable sensor technology. For decades, the OSH paradigm has been predominantly reactive, relying on incident investigations, periodic inspections, and lagging indicators to inform safety improvements. This approach, while valuable, inherently waits for harm to occur before action is taken. Today, digital technologies offer a paradigm shift towards a proactive, predictive, and preventative model of workplace safety.

This potential is not abstract. It manifests in concrete, life-saving applications. AI-powered computer vision systems can analyze live video feeds from a construction site or factory floor to detect, in real-time, when a worker enters a hazardous zone, fails to wear the correct Personal Protective Equipment (PPE), or adopts an ergonomically unsafe posture. Machine learning algorithms can analyze vast datasets of historical incident reports, near-misses, and environmental conditions to predict and flag high-risk scenarios *before* an accident happens, allowing for targeted intervention. Wearable devices, far beyond simple fitness trackers, can monitor a lone worker's vital signs for signs of heat stress or fatigue, detect a fall, or measure exposure to invisible chemical or radiological hazards, triggering an immediate alert. These technologies are the key to moving beyond incremental safety improvements and making a quantum leap towards the EU's goal of "Vision Zero" for work-related deaths.

The Central Thesis

The European Commission's "Digital Package on Simplification" and the associated Digital Omnibus initiative represent a critical and timely opportunity to ensure that this potential is realized. The EU's digital acquis, encompassing landmark regulations like the AI Act, the General Data Protection Regulation (GDPR), the Data Act, and various cybersecurity frameworks, is constructed upon the laudable and essential principles of protecting citizens' rights, privacy, and security. However, from the perspective of OSH professionals, the cumulative effect of these rules has created a web of regulatory complexity that inadvertently stifles the development and deployment of the very technologies designed to uphold the most fundamental right of all: the right to a safe and healthy working environment.

This report's central thesis is that the current digital rulebook, when applied to the unique context of occupational safety, is not yet fit for purpose. It suffers from a series of critical misalignments, ambiguities, and disproportionate burdens that act as significant "bureaucratic hurdles". These frictions impose prohibitive costs on innovative SMEs, create profound legal uncertainty for employers seeking to enhance safety, and establish practical impossibilities in reconciling data protection with the operational necessities of safety-critical

systems. The Digital Omnibus initiative provides the ideal vehicle to "stress-test" these digital rules and make the immediate, targeted adjustments necessary to unlock the life-saving potential of technology in the workplace without compromising the EU's high standards.

Report Structure Overview

This report provides a comprehensive analysis of the regulatory barriers and proposes a concrete pathway for reform. It is structured into three main sections:

- Section 1: The AI Act — Re-calibrating 'High-Risk' for Practical OSH Application. This section performs a deep-dive analysis of the AI Act, arguing that its current 'high-risk' classification is a blunt instrument that fails to differentiate between worker surveillance and worker protection, thereby imposing an unnecessary and counter-productive compliance burden on beneficial OSH technologies.
- Section 2: Navigating the Data Labyrinth — GDPR, Data Act, and Cybersecurity. This section examines the complex interplay of data-related regulations, highlighting the inherent paradox between the GDPR's data minimization principle and the data requirements of effective AI, the technical fallacy of anonymization in real-time safety systems, and the need to streamline fragmented cybersecurity reporting obligations.
- Section 3: A Blueprint for Reform — Actionable Recommendations for the Digital Omnibus. This final section translates the analysis into a set of specific, actionable proposals for legislative and policy changes. It outlines a clear path forward for creating an OSH-focused approach within AI regulatory sandboxes, mandating harmonized standards, providing pragmatic data governance guidance, and unifying incident reporting.

Section 1: The AI Act — Re-calibrating 'High-Risk' for Practical OSH Application

The Regulation on Artificial Intelligence (the AI Act) is a landmark piece of legislation, rightly establishing a risk-based framework to ensure that AI systems deployed in the Union are safe, transparent, and respect fundamental rights. The principle of categorizing AI systems based on risk is sound and necessary. However, its application to the specialized domain of occupational safety and health demonstrates a critical lack of nuance. The Act's current structure, particularly its definition of 'high-risk' systems, is overly broad and creates a series of unintended consequences that disproportionately burden, and thereby discourage, the development of technologies specifically designed to enhance worker safety.

1.1 The 'High-Risk' Classification: A Blunt Instrument for a Nuanced Field

The core of the problem lies in how the AI Act defines and classifies high-risk systems in the context of employment. This broad-brush approach fails to make the crucial distinction

between AI systems used for administrative or disciplinary management and those used exclusively for the protection of workers' health and safety.

The Default Classification

Under Annex III of the Act, AI systems intended for use in "employment, management of workers and access to self-employment" are automatically classified as high-risk. This category explicitly includes systems used to "monitor and evaluate the performance and behavior of persons in work-related relationships". From a functional standpoint, this definition unavoidably captures a vast range of beneficial OSH technologies. For example, an AI system that monitors a driver's gaze to detect fatigue, a computer vision system that tracks ergonomic posture to prevent musculoskeletal disorders, or a wearable sensor that monitors heart rate variability to prevent heat stress are all, technically, "monitoring the behavior" of a worker. By default, these protective systems are swept into the same high-risk category as AI used for performance reviews, task allocation, or disciplinary actions, despite their fundamentally different purpose and impact on workers' rights.

The Derogation Dilemma

Article 6 of the AI Act offers a potential off-ramp, allowing a provider to derogate from the high-risk classification if their system "does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons". While this appears promising, the conditions for applying this derogation are ambiguous and poorly suited to the OSH context. The derogation applies if the AI system is intended to perform a "narrow procedural task," "improve the result of a previously completed human activity," or perform a "preparatory task".

An OSH AI system does not fit these descriptions. A system that provides a real-time alert to a worker about to enter an area with a toxic gas leak is not merely "preparatory"; its entire purpose is to intervene directly to prevent the most significant harm possible. Its function is immediate and preventative, not an improvement on a past action. This legal ambiguity creates a significant dilemma for developers, particularly SMEs with limited legal resources. Faced with the choice between arguing a complex and uncertain legal case for derogation or accepting the default high-risk classification, the rational and risk-averse choice is to default to 'high-risk', thereby triggering the full suite of costly and burdensome compliance obligations.

This situation creates a damaging paradox. A regulatory framework designed with the protection of health and safety as a primary objective inadvertently disincentivizes the very technologies created for that purpose. The legal uncertainty surrounding the derogation, combined with the threat of staggering fines for non-compliance—up to €35 million or 7% of global turnover for certain violations —creates a powerful "innovation chilling" effect. Innovators in the OSH space are forced to either abandon promising projects or channel scarce resources away from research and development into navigating a compliance pathway that was not designed with their products in mind. This ultimately undermines the Act's goal

of protecting EU citizens, as fewer advanced safety tools reach the market, or those that do are prohibitively expensive.

The Prohibited List and OSH Context

The Act's list of prohibited AI practices in Article 5 also highlights this lack of contextual nuance. It rightly bans "emotion recognition in workplaces or educational institutions". This is a crucial protection against intrusive and discriminatory management practices. However, the Act provides an exception for "medical or safety reasons", but fails to define this exception with sufficient clarity. This ambiguity could inadvertently prohibit the development of a system designed to monitor the cognitive state of an airline pilot or a nuclear power plant operator for signs of extreme stress or fatigue—conditions that are direct precursors to catastrophic safety failures. A blanket ban, without a clearly defined and workable safety exception, risks throwing out the baby with the bathwater, preventing legitimate, life-saving applications in the name of preventing misuse in other contexts.

1.2 Deconstructing the Compliance Burden for OSH Technology

Once an OSH AI system is classified as high-risk, it becomes subject to a stringent set of requirements under Chapter 2 of the Act. While these requirements are well-intentioned, their generic nature makes them particularly burdensome and, in some cases, practically impossible to meet in the dynamic and unpredictable context of workplace safety.

Risk Management Systems (Art. 9)

Article 9 requires providers to establish a risk management system that identifies, evaluates, and mitigates all "known and foreseeable risks" to health, safety, and fundamental rights throughout the system's lifecycle. The challenge for OSH is the term "foreseeable." Workplace environments are inherently unpredictable. While many risks are known (e.g., falls, chemical exposure), a key value proposition of advanced AI is its ability to detect novel or emergent risks—the "black swan" events that traditional risk assessments miss. The very purpose of some OSH AI is to identify and prevent the *unforeseeable*. Requiring a provider to document *ex-ante* all foreseeable risks of a system designed to manage the unforeseeable creates a compliance trap. Furthermore, defining "reasonably foreseeable misuse" is exceptionally difficult. Could a worker deliberately trigger a false alarm? Could a system be hacked to create a hazard? While possible, documenting and mitigating every conceivable misuse scenario places an immense and impractical burden on developers.

Data Governance (Art. 10)

Article 10 mandates that high-risk systems be trained, validated, and tested on datasets that are "relevant, sufficiently representative, and to the best extent possible free of errors and complete". This presents a fundamental conundrum for OSH AI. Data on severe and fatal workplace accidents is, thankfully, rare. It is impossible to create a "complete" and "representative" dataset of catastrophic equipment failures or fatal injuries without

deliberately exposing workers to unacceptable risks. Developers must therefore rely on datasets of near-misses, simulated data, or synthetic data. While these are valid and necessary techniques, the "completeness" and "representativeness" of such datasets can always be challenged by a regulator, creating a persistent state of legal uncertainty for the provider. The standard is laudable but practically unattainable for the very scenarios where AI could provide the most value.

Technical Documentation (Art. 11) and Record-Keeping (Art. 12)

The Act requires providers to create and maintain extensive technical documentation as outlined in Annex IV, demonstrating compliance with all requirements. It also mandates that high-risk systems automatically log all events in a tamper-resistant manner to ensure traceability. For SMEs, which form the backbone of the specialized OSH tech industry, the administrative overhead of this level of documentation is a significant resource drain. It diverts critical funds and engineering talent away from core R&D into compliance paperwork. The logging requirement, when applied to a real-time system processing high-frequency data from hundreds of IoT sensors across a worksite, can create enormous technical challenges and costs related to data storage, management, and security, with questionable added value for safety outcomes.

Human Oversight (Art. 14)

Article 14 requires that high-risk systems be designed to ensure "effective human oversight". This is a critical principle. However, the Act offers no concrete guidance on what "effective" means in practice. In a fast-paced industrial environment, who is the designated human overseer? Is it a control room operator, a floor supervisor, or a remote OSH manager? What specific competencies, training, and authority must this person possess? A generic requirement for "human oversight" fails to recognize that meaningful oversight of a complex OSH AI system requires the specific expertise of a qualified and certified OSH professional, not just any available employee. This ambiguity does not just create a compliance challenge for the provider; it creates a significant and undefined liability risk for the employer who deploys the system.

1.3 The Prohibitive Cost of Conformity

The cumulative effect of the high-risk classification and its associated compliance burdens is a financial barrier that is insurmountable for many of the most innovative players in the OSH technology space.

Direct and Indirect Compliance Costs

The direct costs of meeting the high-risk requirements are substantial. Independent analysis, based on the Commission's own impact assessment, estimates that the total compliance costs for a single high-risk AI product can be up to €400,000 for an SME. This could lead to a profit reduction of as much as 40%, a crippling blow for a small or medium-sized enterprise.

The total cost of the AI Act to the European economy is projected to be €31 billion over the next five years. For many high-risk systems, these costs are compounded by the need for mandatory third-party conformity assessments, which add further expense and significant delays to the product lifecycle.

The indirect costs are just as damaging. The combination of high upfront investment, profound legal uncertainty, and a prolonged time-to-market creates a hostile environment for the very SMEs and startups that are the primary engines of innovation in niche, specialized fields like OSH technology. This risks creating a market where only large, generalist technology firms can afford to compete, firms that may lack the deep, domain-specific expertise required to build truly effective OSH solutions. The worst-case scenario is a market vacuum, where potentially life-saving technology is simply not developed or brought to market in the EU due to these regulatory barriers.

This entire dynamic creates a perverse incentive for employers to maintain the status quo of "dumb" safety. Consider an employer facing a persistent safety issue, such as the risk of vehicle-pedestrian collisions in a warehouse. They have two choices. The first is the traditional approach: painting lines on the floor, installing physical barriers, and conducting periodic worker training. This approach is moderately effective, and its regulatory and liability implications are well-understood. The second choice is a new, AI-powered system that uses computer vision and wearable sensors to provide real-time, dynamic warnings to both drivers and pedestrians. This system is potentially far more effective. However, as the "deployer" of this 'high-risk' AI system, the employer faces a host of new, complex, and ill-defined obligations regarding human oversight, data management, and monitoring. The cost of the system is inflated by the provider's own compliance burden. The legal landscape is new, and the potential for severe fines for non-compliance is high. Faced with this choice between a moderately effective, low-regulatory-risk solution and a highly effective but high-cost, high-regulatory-risk alternative, the rational economic decision for many businesses will be to avoid the AI system. The regulation, therefore, in its pursuit of ensuring AI is safe, inadvertently promotes the continued use of less safe, non-AI alternatives, directly contradicting the fundamental goal of improving worker safety.

Section 2: Navigating the Data Labyrinth — GDPR, Data Act, and Cybersecurity

The effectiveness of any advanced OSH technology, particularly those powered by AI, is fundamentally dependent on data. These systems require data to learn, to monitor, and to intervene. However, the regulatory landscape governing data in the EU is a complex patchwork of interconnected and sometimes conflicting requirements. The General Data Protection Regulation (GDPR), the Data Act, and various cybersecurity reporting mandates create a "data labyrinth" that is exceptionally difficult for OSH technology developers and deployers to navigate, posing further barriers to the adoption of life-saving innovations.

2.1 The GDPR Paradox: Data Minimization vs. Predictive Accuracy

At the heart of the data challenge lies a fundamental tension between a core principle of the GDPR and a core requirement of effective AI. This paradox must be resolved if data-driven safety systems are to flourish.

The Principle of Data Minimization

Article 5 of the GDPR enshrines the principle of data minimization, which requires that personal data collected must be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed". This is a cornerstone of modern data protection, designed to prevent the excessive collection and misuse of personal information.

The Needs of OSH AI

In stark contrast, effective and reliable AI systems, especially those based on machine learning, are data-hungry. To train a model that can accurately distinguish between a safe and an unsafe ergonomic posture, or predict the likelihood of an equipment failure based on subtle vibration patterns, requires a large, rich, and comprehensive dataset. A model trained on "minimal" data will be inherently less accurate, less reliable, and more prone to biases that could disproportionately endanger certain groups of workers. This creates a direct and unavoidable conflict: the very data that a safety engineer deems essential for building a reliable, life-saving system may be viewed as excessive by a data protection officer applying a strict interpretation of the data minimization principle. The regulation pits two laudable goals—privacy and safety—against each other.

Legal Basis for Processing Health Data

This tension is amplified when the data in question is health data, which is common in advanced OSH systems (e.g., monitoring for heat stress, fatigue, or exposure). Health data is classified as "special category data" under Article 9 of the GDPR, and its processing is prohibited by default unless a specific exception can be met. For the OSH context, the most relevant legal bases for processing are:

- Legal Obligation (Article 6(1)(c) and 9(2)(b)): This basis applies when processing is necessary for the employer to comply with their legal obligations under national OSH laws, which universally require employers to ensure the health and safety of their workers. This is the most robust and appropriate legal basis for deploying a certified OSH safety system.
- Occupational Medicine (Article 9(2)(h)): This applies when processing is necessary for assessing the working capacity of an employee, for example, through an occupational health assessment.
- Legitimate Interest (Article 6(1)(f)): This basis is often considered but is problematic in the employment context. It requires the employer to conduct a complex balancing test, weighing their interest (e.g., preventing accidents) against the employee's fundamental rights and freedoms. Given the inherent power imbalance in the employer-employee

relationship, reliance on legitimate interest for processing sensitive health data from wearables is legally precarious and generally discouraged.

The core issue is a lack of explicit, EU-level guidance that definitively confirms that an employer's duty to provide a safe workplace under established OSH directives constitutes a sufficient "legal obligation" to justify the processing of necessary health and biometric data via a properly implemented, privacy-preserving, and certified OSH AI system. This ambiguity leaves employers legally exposed and hesitant to adopt such technologies.

2.2 The Anonymization Fallacy in Real-Time Systems

A frequently proposed solution to the GDPR paradox is the anonymization or pseudonymization of data. The logic is that if data is no longer personally identifiable, the strictest requirements of the GDPR do not apply. The Data Act, for instance, explicitly excludes fully anonymized data from its scope. However, for real-time OSH systems, this solution is often a fallacy based on a misunderstanding of both the technology and the operational requirements of safety.

The Technical Impracticality

True, irreversible anonymization of the kind of data streams generated by modern OSH systems—which are often high-frequency, multi-variate (combining location, motion, and biometric data), and longitudinal (tracked over time)—is exceptionally difficult and often impossible without destroying the data's utility. Anonymization techniques like suppression (removing identifiers) or generalization (reducing precision, e.g., replacing an exact location with a broad zone) degrade the quality of the data to the point where it becomes useless for the real-time detection of hazards. An AI model cannot accurately detect an ergonomic risk if it doesn't have precise limb position data, nor can it alert to a man-down situation if it doesn't have a specific location.

The Risk of Re-identification

Even when data is pseudonymized (replacing direct identifiers with a token), the risk of re-identification remains extremely high. In a workplace context, the combination of quasi-identifiers such as a worker's specific role, work area, and shift pattern can often be enough to uniquely identify an individual through linkage attacks. More fundamentally, a safety system must, at the moment of intervention, be able to identify the specific individual who is in danger. A system that can only send an alert saying "an unidentified person is in danger in a general area" is not an effective safety system. It creates a catch-22: to be compliant with a simplistic interpretation of data protection rules, the system must be rendered functionally ineffective.

This entire debate highlights how the current regulatory discourse is lagging behind technological development. The framework remains largely stuck in a binary choice between fully identifiable personal data and fully anonymized data. It fails to adequately recognize or

incentivize a third way: the use of Privacy-Enhancing Technologies (PETs). Solutions like federated learning, where AI models are trained on decentralized data without the raw data ever leaving a local device or server, offer a powerful means of building effective models without centralizing sensitive information. The use of synthetic data generation can create statistically realistic datasets for training and testing without using any real personal data at all. Techniques like differential privacy and secure multi-party computation provide further avenues for analysis while offering strong, mathematically provable privacy guarantees. The current regulations do not provide a clear "safe harbor" or "presumption of conformity" for organizations that invest in and deploy these state-of-the-art PETs. This legislative gap forces innovators and employers back into the unworkable anonymization debate, rather than encouraging them to adopt the very technologies that can reconcile the goals of safety and privacy.

2.3 Streamlining the Cybersecurity Reporting Framework

The final layer of complexity in the data labyrinth is the fragmented and demanding set of rules for reporting cybersecurity incidents. OSH technologies, particularly connected devices like wearables and IoT sensors, fall under the purview of multiple regulations, each with its own reporting triggers and timelines.

A Web of Obligations

An OSH technology provider or the employer deploying the technology could face reporting obligations under several parallel regimes:

- The GDPR requires notification to a supervisory authority of a personal data breach, typically within 72 hours.
- The NIS2 Directive requires operators of essential services to report significant security incidents.
- The Cyber Resilience Act (CRA), which applies to all "products with digital elements," introduces new and extremely stringent reporting obligations.

The 24-Hour Reporting Window

The CRA is particularly concerning from an OSH perspective. It mandates that manufacturers report any *actively exploited vulnerability* to their national Computer Security Incident Response Team (CSIRT) and ENISA within 24 hours of becoming aware of it. While rapid reporting is important for cybersecurity, this extremely short, rigid timeline creates a potential conflict with the overriding priority of worker safety.

In the immediate aftermath of a cybersecurity incident affecting an operational OSH system, the primary focus of all available technical and safety personnel must be on ensuring the safety of the workers on the ground. This involves assessing the impact on safety functions, implementing manual overrides or backup procedures, and communicating directly with the affected workforce. Diverting these key personnel to gather information and file a detailed administrative report to meet a 24-hour deadline can be dangerously counter-productive. It

risks prioritizing bureaucratic compliance over the immediate, hands-on management of a potentially life-threatening situation. The Digital Omnibus's stated goal of simplifying and streamlining incident reporting obligations is therefore not just a matter of reducing administrative burden; from an OSH standpoint, it is a matter of ensuring that regulatory requirements do not interfere with the primary duty of care to protect human life.

Section 3: A Blueprint for Reform — Actionable Recommendations for the Digital Omnibus

The analysis presented in the preceding sections demonstrates a clear and urgent need for targeted reforms to the EU's digital acquis. These reforms are necessary to remove the unintended barriers that are currently stifling the development and adoption of life-saving OSH technologies. This section translates that analysis into a blueprint of five concrete, actionable proposals for the European Commission to consider for inclusion in the Digital Omnibus proposal and the forthcoming Digital Fitness Check. These recommendations are designed not to weaken protections, but to create smarter, more nuanced, and technologically-aware regulations that actively foster safety innovation.

3.1 Proposal: OSH-Focused Streams in AI Regulatory Sandboxes

The AI Act, in Article 57, wisely provides for the establishment of "AI regulatory sandboxes" to foster innovation by allowing for the development, testing, and validation of AI systems in a controlled, supervised environment. This mechanism is the ideal tool for addressing the specific challenges faced by OSH AI technologies.

Recommendation

It is proposed that the Commission, through the Digital Omnibus or subsequent implementing acts, explicitly mandates or strongly encourages Member States to establish dedicated Occupational Safety and Health (OSH) focused streams within their national AI regulatory sandboxes.

These specialized streams would create a "safe harbor" where innovative SMEs and startups can develop and test their technologies without the immediate threat of the severe administrative fines associated with the high-risk classification. Crucially, these sandboxes would serve as a collaborative platform, bringing together technology developers, certified OSH professionals, worker representatives, social partners, and national competent authorities—including representatives from bodies like EU-OSHA. This collaborative environment would allow for the co-development of best practices and provide much-needed clarity on the AI Act's ambiguous requirements in a real-world context. For instance, participants could work together to define what constitutes "effective human oversight" for a specific type of OSH system or establish practical criteria for assessing a "significant risk of harm".

The benefits of this approach are manifold. It would de-risk innovation for SMEs, allowing them to focus on technological excellence rather than legal navigation. It would provide regulators with invaluable insights into the practical application of the AI Act, informing future guidance and evidence-based policymaking. For participants who successfully complete the sandbox program, the resulting documentation and exit report could be used to demonstrate compliance and, as envisioned in the Act, lead to an accelerated conformity assessment process, significantly reducing time-to-market for proven, life-saving products.

3.2 Proposal: Mandate Harmonized Standards and Professional Certification

One of the most effective tools in the EU's single market framework for reducing regulatory complexity is the use of harmonized standards, which provide a "presumption of conformity" with the essential requirements of a regulation. This proven approach should be systematically applied to high-risk AI in the OSH domain.

Recommendation 1: Harmonized Standards for OSH AI

The Commission should issue a formal standardization request to the European Standards Organisations (CEN/CENELEC) to develop a specific set of harmonized standards for AI systems intended for OSH applications. This process must be conducted with the active participation of key stakeholders, including EU-OSHA, national OSH institutes, OSH professional bodies, and social partners. These standards would translate the abstract, high-level requirements of the AI Act's Articles 9 through 15 (covering risk management, data governance, documentation, transparency, human oversight, and robustness) into concrete, auditable technical specifications and processes tailored to the realities of the workplace. For an SME developer, compliance would become a clear engineering objective rather than a vague legal challenge. Demonstrating conformity with the harmonized standard would provide a legal presumption of conformity with the Act itself, offering immense legal certainty and dramatically reducing compliance costs.

Recommendation 2: A "Certified AI Safety Professional" (CAISP)

The AI Act's requirement for "effective human oversight" is critical but undefined. The complexity of modern OSH AI systems requires a level of expertise that goes far beyond that of a typical manager or IT professional. To address this gap, it is proposed that the EU facilitate the creation of a new, recognized professional certification: the "Certified AI Safety Professional" (CAISP).

This certification, developed in line with established professional standards, would ensure that individuals tasked with specifying, deploying, auditing, and overseeing high-risk OSH AI systems possess the necessary interdisciplinary skills in OSH principles, AI technology, data protection, and regulatory compliance. This would provide employers with a clear benchmark for competence and regulators with a tangible measure of "effective oversight." It aligns directly with calls from professional OSH organizations for enhanced training, guidance, and

ethical frameworks for the use of AI in safety.

3.3 Proposal: Pragmatic Guidance on Data Governance for Safety Systems

The current conflict between the GDPR's data minimization principle and the data requirements of effective AI must be resolved with clear, pragmatic guidance from EU authorities.

Recommendation 1: Endorse Privacy-Enhancing Technologies (PETs)

The Commission, in coordination with the European Data Protection Board (EDPB), must issue formal guidance that moves beyond the unworkable binary of "identifiable vs. anonymized" data for real-time safety systems. This guidance should explicitly recognize the use of state-of-the-art Privacy-Enhancing Technologies (PETs) as a valid and preferred method for achieving compliance with both the GDPR's data protection principles and the AI Act's data quality requirements. Specifically, the guidance should confirm that techniques such as federated learning, synthetic data generation, and differential privacy, when properly implemented, can satisfy the principle of "data protection by design and by default". This would create a powerful incentive for industry to invest in and adopt these advanced, privacy-preserving solutions.

Recommendation 2: Clarify the Legal Basis for Processing

To remove the legal ambiguity that currently chills adoption by employers, the EDPB and the AI Office should issue joint guidance that clarifies the legal basis for processing worker health data for safety purposes. This guidance should affirm that an employer's legal obligation under the OSH Framework Directive (89/391/EEC) and its progeny to ensure a safe and healthy workplace constitutes a valid and sufficient legal basis under GDPR Article 6(1)(c) and Article 9(2)(b) for the processing of necessary health and biometric data, provided that this processing is done via a certified, high-risk OSH AI system that adheres to harmonized standards and incorporates PETs. This would provide employers with the legal certainty they need to confidently deploy these life-saving systems.

3.4 Proposal: A Unified and Risk-Based Incident Reporting Mechanism

The Digital Omnibus initiative's goal of streamlining cybersecurity reporting is strongly supported from an OSH perspective, as it addresses the current fragmented and overly burdensome landscape.

Recommendation

The proposed unified reporting mechanism should be designed around a risk-based, tiered reporting timeline. This approach would ensure that the urgency of the reporting obligation is proportional to the severity of the risk to health and safety. A critical distinction must be made

between two tiers of incidents:

- Tier 1: Imminent Safety Threat or Major Personal Data Breach. This tier would cover incidents that directly and immediately endanger the health and safety of workers (e.g., a malicious hack that disables a safety interlock system) or involve a large-scale breach of sensitive worker health data. For these high-severity incidents, a short reporting timeline of 24 to 72 hours is appropriate and necessary.
- Tier 2: System Vulnerability or Minor Incident. This tier would cover the discovery of a software vulnerability with no evidence of active exploitation in a safety-critical environment, or a minor operational incident with no immediate safety impact. For these incidents, a longer reporting timeline (e.g., 7 to 14 days) should be permitted. This would allow the organization to prioritize its resources on investigating, containing, and remediating the issue first, ensuring operational safety is restored before diverting personnel to administrative reporting tasks.

This tiered approach aligns with the core simplification goal of the Digital Omnibus while ensuring that regulatory duties do not inadvertently compromise on-the-ground safety management in the critical moments following an incident.

The following table summarizes these key proposals, mapping the identified regulatory issues to concrete solutions and their intended impact.

Table 1: Summary of Proposed Modifications to the EU Digital Acquis for the Advancement of Occupational Safety and Health

Regulation	Identified Issue	Proposed Modification	Rationale / Intended Impact
AI Act	Overly broad 'high-risk' classification chills innovation in beneficial OSH AI systems.	Issue Commission guidance, co-authored with EU-OSHA, clarifying the Art. 6 derogation criteria for OSH systems to differentiate between surveillance and protective monitoring.	Reduces legal uncertainty; prevents over-classification of life-saving tech; encourages SME investment in OSH AI.

AI Act	Ambiguous and burdensome compliance requirements (Art. 9-15) are disproportionately costly for OSH tech SMEs.	Issue a standardization request to CEN/CENELEC for harmonized standards for OSH AI. Create an EU-level certification for "Certified AI Safety Professionals" (CAISP).	Provides a clear "presumption of conformity," drastically reducing compliance costs and complexity. Ensures genuinely qualified human oversight.
AI Act / GDPR	Lack of a safe, cost-effective pathway for validating innovative OSH AI.	Establish dedicated OSH-focused streams within AI Regulatory Sandboxes.	Accelerates innovation and market access for SMEs by providing a controlled testing environment with regulatory guidance and reduced liability.
GDPR / Data Act	Conflict between data minimization principle and the need for comprehensive data for effective AI safety models. Impracticality of anonymizing real-time safety data.	Issue guidance recognizing Privacy-Enhancing Technologies (PETs) as a primary means of compliance. Clarify that an employer's legal OSH obligation is a valid basis for processing necessary health data.	Resolves the "data vs. privacy" paradox. Moves beyond the unworkable anonymization debate and encourages modern, privacy-preserving technical solutions. Provides legal certainty to employers.

Cyber Resilience Act / NIS2 / GDPR	Fragmented, overlapping, and overly demanding cybersecurity incident reporting timelines that can conflict with immediate on-the-ground safety priorities.	Implement a unified reporting portal with a risk-based, tiered timeline that distinguishes between imminent safety threats and lower-level system vulnerabilities.	Reduces administrative burden. Ensures that operational safety response is always prioritized over premature administrative reporting. Aligns with the Digital Omnibus's core simplification goal.

Conclusion: From Regulatory Burden to Technological Enablement

The recommendations outlined in this report are not a plea for deregulation or a weakening of the EU's high standards for safety and fundamental rights. On the contrary, they are a call for smarter, more precise, and technologically-aware regulation. The current framework, in its laudable effort to be comprehensive, has become a blunt instrument. It fails to make the critical distinction between AI used for invasive worker surveillance and AI used for essential worker protection. The path forward requires nuance: regulation that can differentiate purpose and context, and that is agile enough to recognize and reward the use of privacy-preserving technological solutions.

The European Union's strategic commitment to "Vision Zero" for work-related deaths is a moral and societal imperative. Achieving this ambitious goal in an increasingly complex industrial world is impossible without embracing the most advanced technological tools at our disposal. AI and IoT systems are not merely efficiency tools; they are the next generation of personal protective equipment, engineered safety controls, and proactive risk management systems. They have the potential to see, predict, and prevent hazards in ways that are simply beyond human capability.

The Digital Omnibus initiative provides a pivotal opportunity to make the necessary course corrections. By re-calibrating the AI Act's risk framework, clarifying the data governance landscape, and creating practical pathways for innovation through sandboxes and harmonized standards, the Commission can transform the EU's digital acquis. It can move from a framework that is perceived as a barrier to one that is seen as an enabler of safety. This is not merely an issue of reducing administrative burdens for businesses or enhancing the competitiveness of the EU's tech sector. It is a fundamental matter of occupational safety and health. The ultimate goal is to create a regulatory environment that actively encourages

and accelerates the adoption of technologies that will save lives, prevent occupational diseases, and solidify the European Union's position as the global leader in ensuring safe and healthy workplaces for all.