

From curiosity to confidence

Using AI to transform enterprise-level environment, health and safety



Shape the future
with confidence

The better the question. The better the answer. The better the world works.

This paper draws on EY teams' industry experience, research into practical use cases, client discussions, and EY multidisciplinary experience to provide practical insights for organisations seeking to move from AI curiosity to confidence.

It explores how leaders can take a structured, responsible, and proactive path toward AI-enabled EHS improvement at an enterprise-level; in particular, by moving through three distinct stages while incorporating four key principles - to better help enable capability, trust, and embedment.

Environment, health and safety (EHS) functions are at a critical turning point. The scope, complexity, and scrutiny of EHS risk has expanded significantly over recent years, while the capacity of organisations to oversee that risk has not kept pace.

Increasing regulatory complexity, exponential data growth, and rising expectations from workers, regulators and the community are placing greater demands on EHS functions. At the same time, EHS teams are under pressure to deliver more, with less.

Against this backdrop, artificial intelligence (AI) has entered the EHS conversation with a growing urgency. AI can help organisations make better sense of the vast volumes of unstructured data which is now experienced throughout the work environment. These systems can also support the interpretation of regulatory complexity more efficiently, and support a shift from reactive incident management toward more proactive harm prevention. More specifically, a shift to agentic AI offers a way to support more automated, data-informed decision-making and dynamic workflow redesign, going beyond traditional automation to transform entire business processes.¹

For many leaders, AI provides not only a mechanism for the EHS function to be more efficient, but also a tangible enterprise-level enabler of EHS improvement.

//
EHS performance has not been improving at the expected pace for some time now. While it's certainly not a silver bullet, AI offers a tangible way for us to shift the dial - better protecting our environment and sending more people home safely every day.

Patricio Estevez
Partner, Environment, Health and Safety, Ernst & Young, Australia

¹ [Agentic AI powers enterprise transformation | EY - US](#)

AI use is becoming more common across EHS functions, but a shift to enterprise-level adoption is essential for improvement

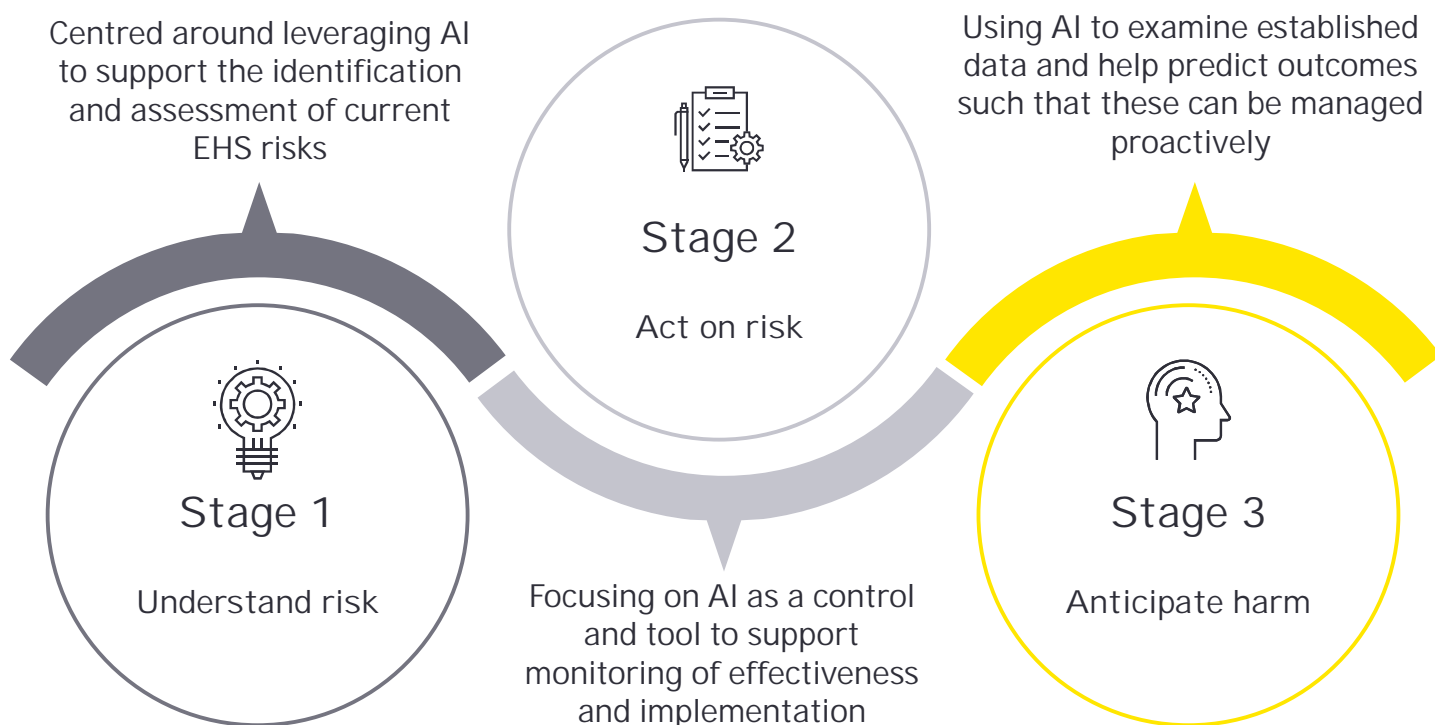
Many organisations are already using AI within their EHS function to improve efficiency. However, this is typically limited to individual application of general AI Chatbots.

While these tools enhance productivity, improving EHS outcomes requires a system-wide approach to risk management and governance. This is because improving EHS is a system-wide risk and governance challenge, not an individual productivity task. Meaningful, organisation-wide EHS improvement depends on enterprise-level AI that can support more integrated oversight, decision-making, and performance uplift.

A practical pathway for AI-enabled EHS improvement across the enterprise

AI-enabled EHS improvement is most logical when it can be clearly linked to effective governance² and, therefore, risk management and broader EHS improvement.

Often, AI roadmaps shift through different types of technology (i.e., machine learning to deep learning, followed by generative AI and agentic AI). However, EY teams see greater benefit in flipping the focus of the AI pathway from the type of AI, to its purpose. This focus will help organisations to build trust and capability in AI as an enabler of EHS improvement:



Key principles which should be adopted across the stages, as well as practical uses of AI at each stage, are discussed further on the following pages.

² [EY architecture for effective health and safety governance](#)

The four principles of AI-enabled EHS improvement

The following principles help address potential barriers to AI enablement and/or EHS improvement, such as organisational and capability constraints, data accuracy risks, governance, privacy, and trust. These principles have been adapted from the [EY Nine Principles of Responsible AI](#), which serve as a foundational framework for organisations developing and deploying AI.

It is critical that these principles are applied across all three stages.

1 | Intention: AI-enabled EHS improvement must begin with a clearly defined purpose. Rather than treating AI as a generic capability or technology upgrade, organisations should explicitly articulate why AI is being used, what EHS problems it is intended to address and what outcomes it is expected to support. This intent should be aligned with the organisation's broader strategic objectives, operating context and risk profile, as well as current and emerging EHS and organisational risks.

It is important to regularly measure and monitor whether the intended purpose of AI is being achieved, and to refine its application as needed to ensure it remains effective and aligned with organisational intent.

2 | Integration: AI only strengthens EHS governance and performance when it is embedded into operating models, decision ownership, data and reporting, and assurance mechanisms. Integration with other operational functions is also critical, considering when data and/or processes could overlap and how to effectively embed this into the design and governance. As with all EHS systems and processes, the needs and capabilities of end users must be considered for the AI to be functional and effective.

3 | Capability: Accurate data, clear ownership, and workforce trust are enabling conditions. Sustained AI adoption in EHS depends on clear accountability across EHS, technology, and risk functions, data that can be trusted and integrated into workflows, and early, transparent engagement with the workforce. While AI can be used to help provide structure to unstructured data (as outlined in Stage 1), if its outputs are found to be incorrect or unhelpful, this presents an opportunity to revisit the data source(s) and adjust.

4 | Human judgement: AI can inform decisions, but accountability for any decisions relating to EHS and associated outcomes, must remain clearly owned by leaders. In safety-critical environments, AI should supplement professional judgement by frontline teams, not replace it. It is important, particularly when it comes to EHS that AI on its own is not used to guide process flow or decision-making.

These principles matter because they determine whether AI strengthens EHS governance and confidence – or simply adds another layer of complexity to an already stretched system.



Stage 1

Understand risk:

Establishing foundational visibility, shifting from reporting to meaningful insight

Many organisations hold vast amounts of EHS data, but much of it remains fragmented, unstructured, or filtered through manual reporting processes – if it is even gathered in the first place. Further, traditional dashboards often present activity (e.g. count of inspections, number of gaps) rather than insight, creating a false sense of control and limiting leaders' ability to discharge due diligence obligations with confidence.³ Effective EHS governance and management requires leaders not only to see, but understand, which risks matter most, where, and why. Time and perspective are required to achieve this, which can often be a capacity and capability challenge for EHS and operational teams.

What is the opportunity presented by AI?

AI can supplement professional judgement, not replace it. It has the potential to provide insights that enable better questions.

This helps shift questions from “Do we have reports?” to “Can we actually see our risk?”. The conversation can then move from discussion of aggregated numbers toward actual risk narratives.

Information relating to EHS risk and performance can also be digested more effectively, with limited reliance on EHS background and capability. Overall, this helps improve the quality of EHS oversight – both from executive and board perspectives, as well as operational leaders who will then need to turn EHS insights into action.

How has AI been used to understand EHS risk?

- Aviation, construction, and industrial safety have used natural language processing (NLP) to classify EHS incidents, extract causal mechanisms, and enable thematic risk analysis at scale.⁴
- AI is being used to continuously integrate and analyse real-time data from sensors, equipment logs, and operational systems to inform safety monitoring dashboards.⁵

What could this look like in practice?

AI can help surface and provide additional context to EHS risk by making existing, but hidden exposure visible. In practice, this can be done in five practical ways:

Standardising free-text EHS inputs into consistent categories

AI can be used to convert free-text (e.g., audit insights, survey responses, toolbox talk conversations, incident descriptions, hazard reports) into consistent, predefined categories (e.g., hazard type, activity, energy source, body part), creating a structured dataset from information that would otherwise remain subjective, inconsistent, and hidden. This helps allow EHS leaders to see volumes, themes and distributions of risk without relying on manual analysis and interpretation.

- Reducing bias in how issues are recorded and surfaced
AI may help reduce variability introduced by different writing styles, experience levels or local practices, enabling more consistency in the visibility of hazards, incidents and controls reported across the organisation. AI can also be used to apply uniform rules when structuring and summarising EHS data, limiting the influence of optimism bias, defensive reporting or inconsistent language, and ensuring that issues are visible even when they are downplayed or described inconsistently. This can be particularly useful for incident investigations and audits or inspections.
- Identifying patterns and drivers of harm
AI can analyse large volumes of EHS and operational data to help identify patterns in emerging risks, recurring causes of hazards and incidents, contributing factors, and systemic issues that are not obvious through manual review. This can support earlier interventions and more informed decision-making across the organisation.
- Linking activity to effectiveness
Once AI has been used to inform understanding of the EHS risk profile (as outlined above), and this has been validated, AI can help monitor ongoing EHS risk by linking controls, actions, and interventions with outcomes. This can help reveal whether controls are effective in practice, i.e., if effort is reducing risk or simply creating activity.
- Prioritising what matters most
AI can help prioritise risks and actions based on severity, exposure and recurrence, so attention and resources are directed where they will reduce harm fastest.

³ More information around this challenge, and what can be done about it more broadly, can be found in an article written by [Roberto Garcia](#) (Partner, EY) in April 2026

⁴ Ricketts, J.; Barry, D.; Guo, W.; Pelham, J. A Scoping Literature Review of Natural Language Processing Application to Safety Occurrence Reports. Safety 2023

⁵ Akano, O. A., Hanson, E., Nwakile, C., & Esiri, A. E. (2024). Designing real-time safety monitoring dashboards for industrial operations: A data-driven approach. *Global Journal of Research in Science and Technology*, 2(2), 1-9

Stage 2

Act on risk:

Embedding insight into control and accountability

Many organisations stall at insight. Risks are known, but action remains inconsistent, reliant on individual effort or informal escalation.

What is the opportunity presented by AI?

AI can be embedded into operational risk management (e.g., to support corrective action management and enable higher order control allocation). Risk insights can be integrated into workflows, planning and prioritisation processes. This helps leaders and managers focus attention where exposure is greatest. AI can also support consistency by reinforcing common thresholds, triggers and escalation pathways across the enterprise.

Through this, controls are not just allocated and monitored; they are implemented and improved.

How has AI been used to act on EHS risk?

- AI analyses real-time air quality data and generates actionable recommendations, such as adjusting traffic flow, modifying industrial operations, or issuing public health alerts when pollution levels exceed safe limits. This enables faster interventions that reduce EHS risk.⁶
- AI-driven computer vision systems have been used for real-time hazard detection linked to automated plant responses, including isolation of machinery when unsafe human-machine proximity is detected. In these situations, AI has been a substitute for manual monitoring and PPE-dependent controls.⁷
- AI-powered drones and robotic systems have been used in confined-space and high-hazard industrial environments specifically to remove workers from exposure. This limits reliance on PPE and procedural controls for hazardous access.⁸

What could this look like in practice?

If *understanding* risk explains where and why harm occurs, *acting* on risk is about using these insights to support consistency, risk-informed decision-making, and effective control implementation.

AI can support action in four practical ways:

- Enable higher-order control allocation
For some hazards (in particular, psychosocial hazards), AI may support elimination, substitution, and/or engineering controls where otherwise only PPE or administrative controls (e.g., training and procedures) were available. In these situations, there is significant potential to improve risk management and control effectiveness which can minimise the potential for harm at the source.
- Improve adoption of EHS processes
AI can help promote awareness and understanding of EHS processes through embedment into EHS Processes. For example, agents can be developed to provide more targeted and concise information to frontline teams regarding the steps they need to follow in high-risk situations.
- Embed risk insights into workflows
AI can help integrate risk signals (e.g., alerts where certain conditions are met and additional steps are required to minimise the potential for harm) into operational processes such as planning, scheduling or approvals. This can help improve the feedback loop and reduce reliance on manual escalation and individual judgement.
- Trigger earlier and more consistent intervention
AI can be used to help detect known threshold breaches or equipment reliability concerns (e.g., visible signs of deterioration). Once detected, AI can prompt or recommend action, supporting both proactive and reactive risk management.
- Strengthen accountability and follow through
AI can support tracking whether actions are taken and help access their effectiveness, reinforcing governance and closing the gap between knowing and doing.

⁶ David B. Olawade, Ojima Z. Wada, Abimbola O. Ige, Bamise I. Egbewole, Adedayo Olojo, Bankole I. Oladapo, Artificial intelligence in environmental monitoring: Advancements, challenges, and future directions, Hygiene and Environmental Health Advances, Volume 12, 2024

⁷ Khurram, M.; Zhang, C.; Muhammad, S.; Kishnani, H.; An, K.; Abeywardena, K.; Chadha, U.; Behdinan, K. Artificial Intelligence in Manufacturing Industry Worker Safety: A New Paradigm for Hazard Prevention and Mitigation. Processes 2025, 13, 1312

⁸ Muthukumar, K., Sundaramahalingam, A., Prathik Jain, S., Sridhar, C.R., Paramaguru, V. (2025). Study on Artificial Intelligence for Industrial Safety. In: Mukesh, R., Venkata Ratnam, D., Theerthamalai, P., Lawrence Raj, P.R. (eds) Proceedings of the 1st International Conference on Advances in Aerospace and Navigation Systems - 2024. iCAAN 2024. Advances in Science, Technology & Innovation. Springer, Cham

Stage 3

Anticipate harm:

Enabling foresight and prevention



The defining characteristic of mature EHS governance is the ability to intervene before harm occurs. However, this is often not possible. While it is known where harm has occurred in the past, leaders often don't have a clear understanding of where it could occur in the future.

What is the opportunity presented by AI?

AI can enable forward-looking insight through predictive modelling, scenario analysis and early warning indicators. Patterns observed over time may help indicate where risk is likely to emerge, degrade or compound. This supports proactive intervention rather than retrospective response. The catch here however is that a solid foundational understanding of risk and the effectiveness of action is needed for this to be possible. Therefore, stages one and two must be implemented before stage three.

This helps governance shift from “What happened?” to “What could cause harm, and what action are we taking now?”. With this, AI can act as a potential enabler of more proactive prevention efforts, allowing EHS professionals to focus on leadership, engagement, and judgement rather than manual analysis.

AI can then become part of the control environment – not just an additional reporting layer.

What could this look like in practice?

If *acting* on risk closes today's gaps, *anticipating* harm is about identifying signals that may indicate risk before it becomes an incident. AI can support anticipation in three practical ways:

- Model how risk accumulates over time to help identify leading indicators of harm
By tracking changes in conditions, controls and exposure, AI can help indicate where risk may be escalating or compounding, even when no event has yet occurred. AI can be configured to notify users that an increase in risk has been detected. Where well established modules exist, AI can be used to analyse early warning signals (e.g., near misses, system deviations or failures, behavioural changes) to help identify patterns that may precede serious (or even catastrophic) events or incidents.
- Surface emerging and non-obvious risks
AI can be particularly useful in exploring complex risks (such as psychosocial harm or contractor fatigue) that do not show up early in traditional lagging indicators. It does this by identifying patterns across diverse data sources to detect early signals, correlations and shifts in behaviour that indicate emerging risk.
- Enable earlier, targeted intervention
Predictive insights from such as emerging risks, early warning signals and control gaps may enable leaders to intervene sooner and more precisely, shifting EHS from reaction to prevention.

How has AI been used to anticipate EHS harm?

- AI analyses large volumes of seismic, satellite, and ground-sensor data to help predict geological hazards (e.g., earthquakes, landslides, volcanic activity), enabling risk mitigation before events occur.⁹
- In advanced manufacturing, AI is used to predict tool failure before it occurs. Machine-learning models analyse real-time data from force, vibration and acoustic sensors to predict cutting tool wear. While this AI has been designed primarily to improve operational efficiency, it also enables tools to be replaced before wear leads to unsafe operating conditions.¹⁰
- AI identifies illegal waste dumping before it causes EHS impacts. AI systems analyse images from drones or surveillance cameras using deep-learning models to locate illegal dumping, enabling greater efficiency over manual methods and allowing authorities to intervene early and prevent soil, water, and ecosystem contamination.¹¹

⁹ Tianjie Zhao, Sheng Wang, Chaojun Ouyang, et al., Artificial intelligence for geoscience: Progress, challenges, and perspectives, The Innovation, Volume 5, Issue 5, 2024

¹⁰ Jones, T., Cao, Y. Tool wear prediction based on multisensor data fusion and machine learning. Int J Adv Manuf Technol 137, 5213-5225 (2025) <https://doi.org/10.1007/s00170-025-15472-4>

¹¹ Fang, B., Yu, J., Chen, Z. et al. Artificial intelligence for waste management in smart cities: a review. Environ Chem Lett 21, 1959-1989 (2023)

So, what now?

The question facing EHS leaders is no longer whether AI has a role to play.

It is how to begin using AI at an enterprise-level responsibly, credibly, and without creating new risk.

With many organisations already encouraging individuals' use of general AI Chatbots, enterprise-level AI is the next step to achieve better EHS outcomes across organisations and industry.



Organisations should begin by using AI to strengthen visibility and understanding of EHS risk.

For organisations early in their AI journey, this means focusing on making existing EHS data usable and reducing blind spots that undermine confidence and due diligence.

For organisations already experimenting with AI, it means reassessing whether current EHS initiatives are genuinely improving EHS outcomes. This is done by considering the aspects such as operating models, risk profiles, workflows, management systems and the needs of end users. A roadmap can then be developed to identify how AI can be leveraged as an enabler of enterprise-wide EHS improvement.

From there, progress should be deliberate. Confidence must be built before scale. Insight must precede action. Action must be embedded before prediction is pursued. At every stage, there must be clear intention, integration with existing ways of working, focus on capability and human judgement embedded in the decision-making process. If these principles are adopted, the benefits of each stage can be better achieved:

Stage 1 Understand risk

- At stage one (Understand risk), AI can help organisations better understand EHS risk by providing insights into why harm may occur, where it is building, and which controls work.

Stage 2 Act on risk

- At stage two (Act on risk), AI can help organisations act on EHS risk by supporting the translation of insight into prioritised, embedded and timely action – without removing human accountability.

Stage 3 Anticipate harm

- At stage three (Anticipate harm), AI can support anticipation of EHS harm by identifying early warning signals and emerging patterns that indicate where future incidents are most likely – before people are hurt.

Organisations that take this approach will move forward steadily. This is not because they are technologically advanced, but because they are clear about what matters, honest about their readiness, and intentional about how AI is used to support safer work.



How can EY help?

EY teams combine their EHS, digital, and AI experience to provide you with the insight you need, no matter where you are in your AI journey. We can support you to:

- Understand how you can make existing EHS data usable.
- Assess whether current AI initiatives are improving EHS governance or adding complexity.
- Assess alignment between current EHS initiatives and desired outcomes, mapping issues, pain points, and deviations from planned workflows.

- Develop a proof of concept and/or business case for a specific enterprise-level AI enabler for EHS improvement.
- Develop a clear and practical roadmap to leverage AI for EHS improvement, based on the readiness of your organisation.

To learn more, reach out to us directly.

Contact us



Karen Mealmaker
Partner, Environment, Health and Safety,
Ernst & Young, Australia
karen.mealmaker@au.ey.com



Roberto Garcia
Partner, Environment, Health and Safety,
Ernst & Young, Australia
roberto.garcia@au.ey.com



Rebecca Dabbs
Partner, Environment, Health and Safety,
Ernst & Young, Australia
rebecca.dabbs@au.ey.com



Patricio Estevez
Partner, Environment, Health and Safety,
Ernst & Young, Australia
patricio.estevez@au.ey.com

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

© 2026 Ernst & Young, Australia
All Rights Reserved.

ED none
SCORE 111888-26-AUNZ

This communication provides general information which is current at the time of production. The information contained in this communication does not constitute advice and should not be relied on as such. Professional advice should be sought prior to any action being taken in reliance on any of the information. Ernst & Young disclaims all responsibility and liability (including, without limitation, for any direct or indirect or consequential costs, loss or damage or loss of profits) arising from anything done or omitted to be done by any party in reliance, whether wholly or partially, on any of the information. Any party that relies on the information does so at its own risk. Liability limited by a scheme approved under Professional Standards Legislation.

ey.com