Digital co-workers in hazardous environments: th

The maintenance sector stands at a critical juncture. Workforce challenges continue to dominate concerns in industrial businesses. In Ultimo's Maintenance Trend Report, the ageing workforce is identified by 63% of respondents as the most pressing challenge impacting maintenance strategy, highlighting the urgency of knowledge transfer and workforce planning. Half (50%) of respondents said recruiting experienced staff was the primary source of disruption over the past year. The industry faces an unprecedented crisis. In hazardous environments governed by DSEAR, ATEX, and IEC regulations, this skills shortage isn't merely an operational inconvenience; it's a safety imperative.

The connection between maintenance and safety has never been more critical. When experienced technicians who understand the nuances of explosive atmospheres and hazardous area compliance retire, they take with them decades of institutional knowledge about what makes equipment safe to operate. This knowledge gap directly impacts environmental, health and safety (EHS) outcomes, creating blind spots that can lead to incidents, regulatory non-compliance, and ultimately, catastrophic failures.

As organisations grapple with these challenges, a new technological paradigm is emerging: agentic artificial intelligence (AI). Unlike traditional automation or even predictive analytics, agentic AI represents a fundamental shift from passive tools to active digital co-workers that can autonomously manage complex EHS workflows whilst maintaining the human oversight essential for safety-critical operations.

From tools to teammates: the agentic revolution

Traditional maintenance technology operates on a reactive basis. Even sophisticated systems require human input for decision-making, with automation limited to rigid, predefined rules. A failure can trigger a work order, a sensor reading possible generates an alert, but the system remains fundamentally passive.

Agentic AI transforms this dynamic entirely. These digital agents operate autonomously within defined parameters, making autonomous decisions, initiating workflows, and learning from every interaction. They don't simply process data; they understand context, identify patterns, and take proactive action. In hazardous environments, this capability becomes transformational.

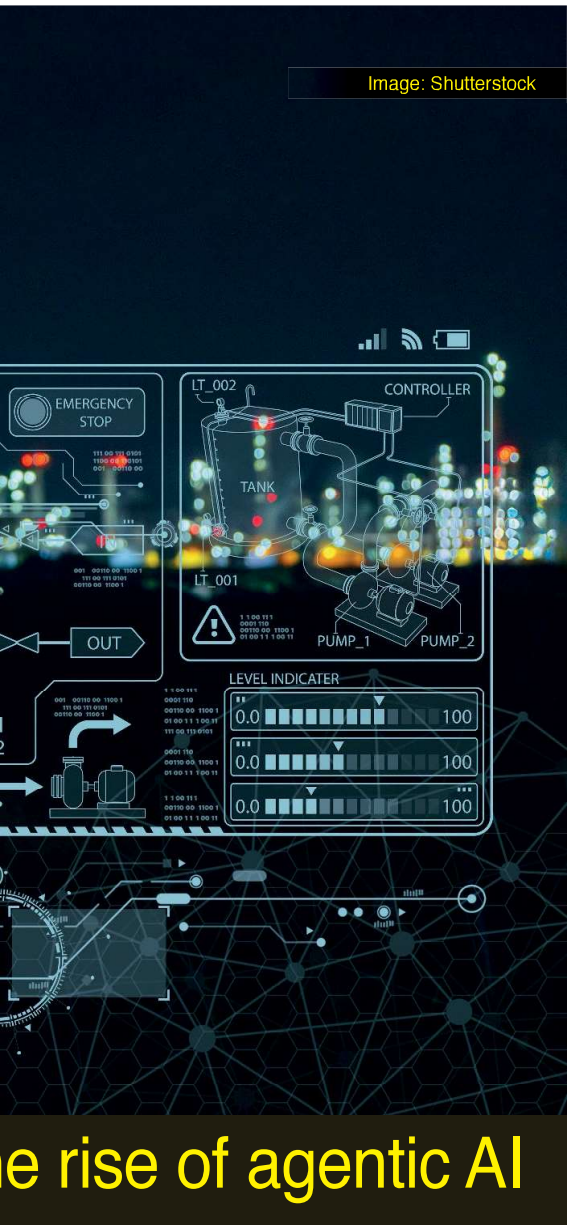

Image: Shutterstock

e rise of agentic AI

Consider the complexity of maintaining equipment in ATEX-certified zones. Traditional approaches rely heavily on human expertise to identify potential safety issues, ensure compliance documentation is complete, and recognise when equipment behaviour suggests a developing hazardous condition. Agentic AI is not bound by the confines of human memory or compliance. It is always active, continuously monitoring work orders, shift logs, and operational data to identify safety patterns that might escape human attention, particularly during periods of high workload or staffing shortages.

Building trust through transparency

The success of agentic AI in hazardous environments hinges on trust, and trust must be earned through transparency and proven reliability. The most effective approach treats the AI agent like an apprentice rather than a replacement. Just as you wouldn't place a new apprentice in charge of critical equipment in an explosive atmosphere, agentic AI begins with simple, supervised tasks before gradually taking on more complex responsibilities.

This apprenticeship model aligns perfectly with safety culture. In hazardous environments, every decision has potential consequences, and the path to autonomy must be carefully managed. The AI agent starts by flagging potential issues for human review, suggesting maintenance schedules, and identifying compliance gaps. Over time, as confidence builds, the agent can take on more autonomous functions whilst maintaining clear audit trails and human oversight capabilities. The AI agent does all the heavy lifting, but the human is always in the loop, and always in the lead.

Transparency remains paramount. Every decision made by the agent must be explainable, every action traceable, and every recommendation based on clear logic that safety professionals can validate. This isn't just good practice; in regulated environments, it's essential for demonstrating due diligence and maintaining compliance with safety standards.

Addressing the under-reporting crisis

One of the most significant challenges in EHS management is the systematic under-reporting of safety incidents, near misses, and unsafe conditions. Traditional safety management systems rely heavily on manual reporting, which creates substantial blind spots. Employees may forget to document events, particularly minor incidents, or may be reluctant to report issues that could reflect poorly on their performance.

This under-reporting problem is particularly acute in hazardous environments where the consequences of missed incidents can be severe. A minor equipment anomaly that goes unreported today could indicate a developing condition that leads to an explosive atmosphere tomorrow.

Agentic AI addresses this challenge by continuously scanning work orders, shift logs, and operational data for indicators of safety incidents. The system can identify patterns that suggest unreported events: repeated repairs on the same equipment, certain types of damage that typically result from specific incidents, or operational anomalies that correlate with safety conditions.

For example, if work orders consistently reference damage to doors in specific locations, the agent can recognise this pattern and correlate it with historical incident data to identify likely unreported vehicle collisions. The agent then automatically generates a comprehensive incident report for human review and validation, giving safety teams unprecedented visibility into workplace risks, enabling them to take adequate preventive action.

Transforming maintenance workflows

The implementation of agentic AI fundamentally changes how maintenance teams interact with technology. The Maintenance Trend Report also reveals that interest in contextual intelligence has surged by 750% since 2023, with 68% of organisations now recognising its potential to positively impact their maintenance practices. Rather than requiring technicians to navigate complex software interfaces or remember to log into multiple systems, the AI agent proactively delivers information when and where it's needed.

Plant managers no longer need to rely only on dashboards to understand the status of their operations. Instead, they can tap into their network of digital coworkers through their preferred communication channels and instantly receive the latest updates. For example, "No issues encountered during the night shift last night" or "Two minor conveyor issues addressed by night shift operators. Recommend bringing in two technicians on overtime to prevent day shift disruption. Cost: £800. Do you want to approve?"

This multimodal approach ensures that managers stay informed in the way that best fits their context and the way they

Image: Shutterstock

work. Rather than discovering problems after they've escalated, managers receive timely information that enables preventive action. This is especially valuable in hazardous environments where situational awareness is critical.

Human-AI collaboration

The ultimate vision for agentic AI in maintenance isn't about replacing human expertise; it's about amplifying it. Every role in the maintenance ecosystem – from planning to execution to compliance management will be in charge of their own ecosystem of specialised digital co-workers. Within the ecosystem, the agents handle routine tasks and provide expert insights. They also coordinate activities: with their human colleague, but also amongst one another.

EHS agents will continuously monitor compliance requirements, ensuring that maintenance activities align with standards such as DSEAR, ATEX, or IEC. Reliability engineering agents will analyse failure patterns and optimise maintenance schedules. Warehousing agents will manage inventory requirements and coordinate parts availability. These agents work together, communicating and coordinating actions to create a more efficient and safer operation.

This collaborative approach addresses the skills shortage by extending the capability of existing teams rather than simply replacing them. A maintenance planner working with an AI agent can handle more complex scheduling, consider more variables, and make more informed decisions than would be possible working alone. The agent handles data processing, pattern recognition, and routine coordination, whilst the human focuses on strategic thinking, complex problem-solving, and safety-critical decisions.

Starting small, scaling smart

The path to successful agentic AI implementation doesn't require wholesale transformation. The most effective approach begins with a single, well-defined workflow where the agent can demonstrate clear value, and build trust. EHS incident reporting provides an ideal starting point because it addresses a recognised problem (under-reporting) with measurable outcomes (increased incident identification and improved safety visibility).

Starting small allows organisations to build confidence, refine processes, and develop the cultural change necessary for broader adoption. Each successful implementation creates a foundation for expanding into additional workflows, gradually building a comprehensive digital workforce that supports human teams across all aspects of maintenance management.

The key is ensuring that each step builds trust through transparency, delivers measurable value, and maintains the human oversight essential for safety-critical operations. Success creates momentum for broader transformation whilst establishing the governance frameworks necessary for managing AI in regulated environments.

Data quality: the foundation of success

One persistent concern about AI adoption is data quality. Many organisations worry that their maintenance records aren't sufficient to support advanced AI applications. However, the most effective agentic AI systems are designed to improve data quality as they operate.

By prompting technicians for additional detail, using computer vision for accurate readings, and applying natural language processing to interpret narrative reports, these systems gradually enhance the availability, quality, and consistency of maintenance data. Better data doesn't just improve AI performance; it enables more sophisticated reliability engineering, more accurate compliance reporting, and more effective safety management. It is at the root of confident, data-driven decision-making.

This improvement in data quality creates a positive feedback loop. Better data

Image: Ultimo

enables more accurate AI insights, which drive better maintenance decisions, which generate higher quality data. Over time, this cycle transforms not just the technology infrastructure but the entire maintenance operation. Data is foundational not only to the success of today, but also the innovation of tomorrow. Most future AI developments will rely heavily on data. This makes the business case for this type of agentic AI two-pronged: as you invest in realising short-term value, you are generating long-term potential.

Regulatory compliance and safety standards

In hazardous environments, regulatory compliance isn't optional, and any technological solution must demonstrate clear alignment with safety standards. Agentic AI systems designed for these environments must incorporate deep understanding of standards and regulations, such as DSEAR, ATEX, and IEC requirements.

The AI agent becomes a compliance multiplier, ensuring that every maintenance activity considers relevant safety standards, that documentation meets regulatory requirements, and that potential compliance gaps are identified before they become violations. Rather than replacing human judgment about safety matters, the agent ensures that safety considerations are consistently applied and properly documented.

It enables organisations to maintain high safety standards even as experienced personnel retire and newer team members develop their expertise. This systematic approach to compliance reduces the administrative burden on skilled technicians whilst improving the consistency and comprehensiveness of safety management. Less admin work means that highly skilled professionals can apply their expertise where it matters most, to the benefit of the assets and the budget that supports them.

The path forward

The industrial maintenance sector stands at the threshold of a fundamental transformation. Agentic AI offers a path to address the convergent challenges of workforce shortages, increasing regulatory complexity, and growing safety demands. However, success requires thoughtful implementation that prioritises human expertise, builds trust through transparency, and maintains unwavering focus on safety outcomes.

The organisations that embrace this transformation whilst maintaining their commitment to a strong safety culture, will gain a significant competitive advantage. They'll have more capable teams, more reliable operations, and more comprehensive safety management. Most importantly, they'll have digital co-workers that enhance rather than replace human expertise, creating workplaces that are both more efficient and safer.

As we move forward, the question isn't whether agentic AI will transform maintenance operations - it's whether organisations will embrace this transformation proactively or find themselves struggling to catch up with more innovative competitors. The technology is ready. Now the real test is whether we are ready to work alongside our new digital colleagues. ■

About the author

Berend Booms has been an Enterprise Asset Management (EAM) Specialist at Ultimo since 2017, operating from the company's Netherlands headquarters. As the architect of Ultimo's global thought leadership and content strategy, Berend bridges the critical gap between detailed product capabilities and measurable business outcomes. His international academic background includes degrees from the University of Oxford, Université Paris 1 Panthéon-Sorbonne, and Leiden University. Berend previously worked at Cisco Systems as Distribution Change Manager, Business Analyst, and Run-the-Business Training Lead.