



# MAINSTREAM

CELEBRATING 30 YEARS

## Industrial Asset Management

# AI REPORT

What 715 Maintenance and Reliability Professionals  
Told Us About AI, Data, Skills, and the Path Forward

Research Partners



# Executive Summary

Artificial intelligence has moved squarely into the vocabulary of maintenance and reliability professionals across Australia and New Zealand. Conference agendas are dominated by it. Vendor demonstrations lean heavily on it. And leadership teams are asking operational managers what their AI strategy looks like, often before those managers have had a chance to work out what the question even means in the context of a crusher, a conveyor, or a fleet of haul trucks.

To cut through the noise, MAINSTREAM conducted a dedicated research program in early 2026, combining a structured survey of 715 asset management and maintenance professionals with a facilitated industry roundtable featuring practitioners from mining, rail, infrastructure, utilities, manufacturing, oil & gas, and facilities asset management. The findings were then set alongside published research from McKinsey, Deloitte, IBM, MIT Sloan Management Review, and Reliable Magazine, as well as insights from the Engineers Australia AI-in-Practice Report (n=2,197, January 2025), thirteen 1-on-1 interviews with senior Maintenance and Reliability leaders, and a specialist masterclass on AI adoption strategy conducted within the MAINSTREAM community.

**The short answer to whether AI will transform asset management in this region: it will, but not for the reasons most vendors claim it will. The longer answer is more nuanced, and more useful.**

Most organisations in the region sit in the early-to-middle stages of AI adoption. More than a third of survey respondents described their organisation as still exploring, with no deployment in place. A further 22 per cent are deploying in limited operational areas, and 22 per cent are running pilot projects. Only around 11

per cent reported that AI is embedded across multiple operational areas. The technology that has gained the widest traction is not, as many assume, predictive maintenance. It is generative AI. More than half of respondents indicated they are already using tools such as ChatGPT, Copilot, Claude, or Gemini for knowledge management tasks, report generation, and everyday productivity. Predictive maintenance and condition monitoring, while still the headline use case in marketing literature, ranked second at 38 per cent.

When it comes to barriers, data quality and poor data foundations topped the list, consistent with every major international study reviewed. Skills shortages, legacy system integration, and unclear return on investment followed close behind.

The Engineers Australia research found that while 72 per cent of engineers agree that generative AI increases productivity, 78 per cent report that learning is self-directed because organisations are not providing formal training — a critical readiness gap. Workforce AI skills and literacy scored just 2.69 out of 5, the lowest of any readiness dimension.

The concerns expressed by practitioners are pragmatic rather than theoretical. The worry is not that AI will replace jobs. It is that younger engineers



will over-rely on AI outputs without the experience to know when those outputs are wrong, that vendor promises will outstrip operational reality, and that accountability will be unclear when an AI-informed recommendation leads to a poor outcome. There is also growing awareness of ‘AI washing’ — the risk that cosmetic applications of AI technology will distract from genuine organisational challenges. These are the concerns of people who operate in environments where getting it wrong has consequences measured in safety incidents, production losses, and regulatory exposure.

Emerging technologies such as edge intelligence and federated learning show promise in addressing data sovereignty and latency challenges critical to the ANZ region, but adoption remains embryonic. More importantly, the

research reveals a fundamental strategic gap: most organisations are applying AI to process bottlenecks for marginal efficiency gains (planning, scheduling, turnarounds), when the real value lies in reinventing asset strategy from static to dynamic — a once-in-a-career opportunity that requires fundamentally different thinking about how assets are deployed and managed.

**The findings point to a clear conclusion. AI is not failing because the technology is immature. It is stalling because the conditions for it to succeed, including clean data, integrated systems, skilled people, and a culture of trust, are not yet in place in most organisations. The report that follows explores each of these dimensions in detail, grounded in the data and in the words of the practitioners navigating this transition every day.**

# About This Research

This report draws on four complementary research streams designed to capture both the breadth and the texture of current AI adoption in asset management and maintenance across Australia and New Zealand.

## The MAINSTREAM AI Survey 2026

A structured online survey was distributed in early 2026 to maintenance, reliability, and asset management professionals across the region. A total of 715 valid responses were collected. The survey covered current AI adoption maturity, applications in use or under evaluation, barriers to adoption, organisational readiness across six dimensions, confidence in future AI value delivery, key concerns, and perceived opportunities.

The respondent base was weighted towards operational practitioners rather than technology vendors or consultants. Around 26 per cent identified as Maintenance Managers, Engineering Managers, or Superintendents. A further 23 per cent held specialist lead roles in planning, reliability, shutdown, or data. Maintenance and reliability engineers made up 18 per cent, with the balance drawn from heads of asset management, planners and supervisors, consultants, and solution providers.

## Industry Roundtable and Masterclass

A facilitated roundtable discussion was conducted as part of MAINSTREAM's 2026 Think Tank series, bringing together practitioners from organisations spanning mining, rail, road and tunnel infrastructure, utilities, healthcare asset management, and several coal and power asset operators. The session focused specifically on AI integration and adoption barriers and was facilitated by a global asset management leader. The roundtable generated qualitative insights that complement and en-

rich the survey data, and direct observations from participants are referenced throughout this report.

In addition, MAINSTREAM conducted a specialist masterclass within its community featuring a global asset management leader delivering deep-dive insights on AI adoption strategy, value barriers, and the opportunities in multi-agent orchestration and dynamic asset strategy. This session revealed critical thinking about where organisations are applying AI (to process bottlenecks for marginal gains) versus where the real value lies (reinventing asset strategy itself).

## Supporting Research and Practitioner Insights

The analysis was further informed by published research from Deloitte (on AI-enabled predictive maintenance), IBM (on the evolution of maintenance from reactive to predictive), MIT Sloan Management Review (on where traditional maintenance approaches fall short), and a detailed article from Reliable Magazine (on the role of human expertise alongside AI in physical asset management).

New sources integrated into this research include the Engineers Australia AI-in-Practice Report (n=2,197, January 2025), which provides national context on AI adoption across the engineering workforce and highlights the generational divide and the self-directed learning phenomenon. These sources provided both statistical benchmarks and conceptual frameworks against which the local findings were tested.



## Methodology and Acknowledgements

The comprehensive content presented in this report was meticulously designed, developed, and authored by the MAINSTREAM research team, drawing upon extensive expertise in market analysis and industry insights across Australia and New Zealand. As a leading research authority in both the ANZ and UK regions, MAINSTREAM brings three decades of collective experience and methodological rigour to this analysis. The authors acknowledge Octave and Accruent for their invaluable partnership and support throughout this research initiative.

## About MAINSTREAM

Founded in 1996, MAINSTREAM is an award-winning B2B community serving asset-intensive industries with research, information, events, training courses, and digital communication solutions that celebrate the successes, accelerate the careers, and optimise the performance of Asset, Reliability, and Maintenance professionals.

# The Current State of AI Adoption

The survey asked respondents to place their organisation on a six-point maturity scale, ranging from “Not started” through to “Optimising.” The results paint a picture of a region that is interested, active, but still predominantly in early stages.

When these figures are grouped, nearly 45 per cent of respondents fall into the “exploring or not started” category, a further 44 per cent are in pilot or early implementation, and only around 11 per cent have reached the point where AI is genuinely scaled or optimised. This distribution is consistent with the broader global picture. National research cited in the MAINSTREAM Asset Management report found that while 72 per cent of asset-intensive organisations are exploring AI applications, more than three-quarters fail to achieve expected returns.

The Engineers Australia report provides important context on the generational dimension of adoption. Across the surveyed engineering workforce, 43 per cent of practising engineers currently use generative AI in their work, compared to 74 per cent of engineering students. This gap signals both an opportunity and a challenge. Younger cohorts are already comfortable with AI tools, but organisations need to ensure this confidence is grounded in actual competence and that experience is not displaced. The same research found that 73 per cent of engineers agree that AI-related skills will be essential in their roles within the next five years, yet 78 per cent report their learning has been self-directed because organisations are not providing structured training.

The masterclass discussion revealed another critical statistic: of organisations surveyed, two-thirds of maintenance teams planned to adopt AI in 2026, but only one-third had successfully implemented pilots. This represents a significant gap between intention and execution.

What stood out from the roundtable was how universal this experience felt across very different industries. A rail infrastructure manager captured the mood concisely:

**We think we’re going okay, but I pretty much think we’re probably touching the top 5% of possibility. And we don’t know what we don’t know.**

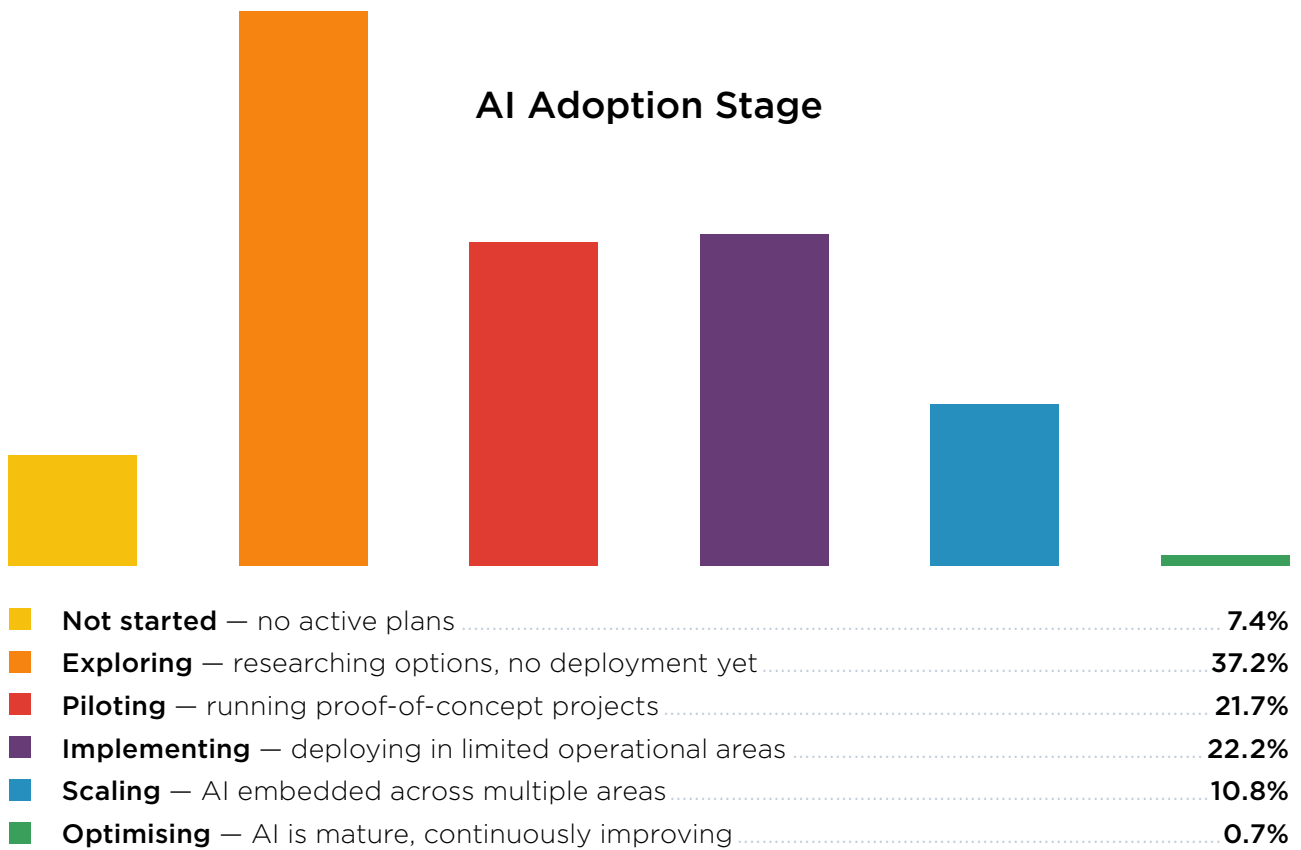
**Rail Infrastructure Manager**

An asset management strategist from a utilities business talked about using generative AI to eliminate routine work but acknowledged this was still the tip of the iceberg. A civil engineer at a power station described AI adoption in their area as next to zero, despite being able to see opportunities clearly. Even a representative from a major global mining company with significant investment in digital strategy, framed their journey as still in its early chapters, with dedicated teams and top-down funding only recently put in place.

The picture is not one of resistance. It is one of organisations at different points on a learning curve, trying to work out where AI delivers real operational value versus where it remains a promising idea that is difficult to execute.



### AI Adoption Stage



# Where AI Is Being Applied

The dominance of generative AI is the standout finding here, and it is worth pausing on. Predictive maintenance has been the flagship AI use case in asset management marketing for years. It remains important, but what has actually gained the most traction on the ground is something quite different: general-purpose language models being used for everyday knowledge work.

Roundtable participants confirmed this pattern consistently. One asset management leader described the impact in practical terms:

**I've been able to churn out outputs that would have taken me days, with all the manual intervention that you normally would have had to have done. Just get rid of the hack work, and then get down to analysing what the result is.**

**Head of Asset Management Strategist,  
Utilities**

Several participants noted their organisations had deployed internal versions of generative AI tools so that employees could query internal documentation without the security concerns of using public models. The value being realised is not predictive in nature; it is administrative and analytical.

Predictive maintenance, at 38 per cent, remains the second most common application, but the roundtable discussion revealed an important qualification. When the facilitator asked for a show of hands on who had predictive analytics actually deployed, only two or three participants out of eight could raise their hand. There is a meaningful gap between “evaluating” predictive maintenance and having it running in production.

A critical insight from the masterclass addresses where organisations typically apply AI. Most commonly, they target process bottlenecks — planning, scheduling, turnarounds — for marginal efficiency gains of 5 to 10 per cent. This is valuable but incremental. The real opportunity lies not in optimising existing static strategies but in reinventing asset strategy itself to be dynamic, responsive, and tied to actual operating goals rather than budgets. This represents a fundamentally different level of value but requires rethinking the entire operating model.


A caution worth noting: practitioners at the masterclass flagged ‘agent washing’ — vendors and organisations labelling routine automation as ‘intelligent agents’ when they are simply process optimisation tools. This mirrors earlier concerns about AI washing more broadly. The distinction matters for setting expectations and avoiding disappointment.

Visual inspection and image analysis, at 22.5 per cent, emerged as a practical entry point that several participants singled out. The WA report highlighted an organisation that built an image recognition tool to identify rolling contact defects from thousands of inspection photos, reducing a task that had previously taken weeks to just a couple of hours. Drone-based inspection, fed into AI platforms, was also cited as a proven application, particularly for structural integrity assessments and thickness testing in hard-to-reach locations.

Digital twins and asset lifecycle planning, while commanding significant attention at conferences and in vendor literature, sat at the bottom of actual adoption. This suggests that these applications, while conceptually appealing, face higher barriers to practical deployment in most organisations.

## AI Application

Respondents were asked which AI applications they are currently using or actively evaluating. They could select as many as applied from a list of nine categories, plus a “none of the above” option.



<span style="color: #0070C0;">■</span> Generative AI (ChatGPT, Copilot, Claude, or Gemini) for knowledge management .....	<b>52.5%</b>
<span style="color: #00BFC4;">■</span> Predictive maintenance / condition monitoring / anomaly detection .....	<b>37.5%</b>
<span style="color: #008000;">■</span> Visual inspection / image and video analysis .....	<b>22.5%</b>
<span style="color: #70AD47;">■</span> Work order / planning optimisation.....	<b>21.7%</b>
<span style="color: #FFD700;">■</span> AI Agents for task and process automation.....	<b>20.8%</b>
<span style="color: #FF8C00;">■</span> Field troubleshooting / root cause analysis support.....	<b>20.0%</b>
<span style="color: #FF0000;">■</span> Drone-based inspection.....	<b>16.7%</b>
<span style="color: #800080;">■</span> Digital twins .....	<b>12.5%</b>
<span style="color: #483D8B;">■</span> Asset lifecycle and investment planning .....	<b>10.0%</b>
<span style="color: #191970;">■</span> None of the above .....	<b>6.7%</b>

# Barriers to Adoption

The survey asked respondents to identify the biggest barriers to AI adoption in their organisation, selecting from a list that covered technical, organisational, and strategic factors. The results confirm what practitioners have been saying in corridors and workshops for the past two years: the obstacles are not primarily about the technology.

## Data Quality and Poor Data Foundations

Data quality was the most frequently selected barrier, chosen by 16 per cent of respondents as a top-five concern. This is consistent with every major study reviewed in preparing this report. MIT Sloan Management Review identifies poor data quality as one of the biggest challenges AI projects face across industries. The MAINSTREAM roundtable participants were even more direct:

**AI amplifies everything – including your data quality problems. We thought intelligent algorithms would compensate for poor data entry, but instead they made our existing issues more visible and more expensive. Garbage in, garbage out applies even more to AI than traditional systems.**

**Maintenance Planning Lead, Manufacturing**

This observation aligns with a critical insight from the CEO of a leading Maintenance AI platform provider: ‘AI won’t fix broken fundamentals.’ The argument is that organisations should treat AI like a reliability program, not a science project. The corollary is stark: if your computerised maintenance management system (CMMS) data is garbage, AI will automate bad decisions faster, not fix them. Start with data fundamen-

tals, then standardise, then layer on AI. Trust beats sophistication.

What made this finding particularly sharp in the local context was the roundtable discussion. An engineering manager managing aged infrastructure at a power station described data collection processes that remain entirely manual, with inconsistent formats accumulated over many years. Even if they wanted to feed an AI system, the effort required to prepare the data would exceed the effort of simply doing the analysis manually. An asset management strategist provided a vivid example:

**The data was a mess. Multiple pieces of information in the one cell in Excel, different date ranges, locations named five different ways. I tried to use Copilot to clean that up, and it could not just do it, because there were just so many different issues going on.**

**Asset Management Strategist, Utilities**

A maintenance manager from a major road and tunnel operator described a similar challenge: multiple asset management platforms that capture data but are not integrated with AI tools, meaning significant work is needed before data can be converted into meaningful insight.

The readiness score for data quality and availability was 2.85 out of 5 across the respondent base. This is a middling score at best, and for many organisations, it reflects years of underinvestment in data governance, master data management, and consistent data capture processes. The roundtable facilitator framed this clearly:



**Data is a process, not a project. You don't clean data and then go, right, I'm done, and then walk away and forget about it, because it doesn't take care of itself. We're changing assets all the time, conditions change all the time. We have to recognise that our processes need to support data capture as part of normal operations.**

**Industry Roundtable Facilitator**

### **Skills Shortages and Digital Literacy**

The lack of in-house AI and digital skills was the second most cited barrier at 12 per cent, and workforce AI skills and literacy received the lowest readiness score of any dimension at 2.69 out of 5. This finding deserves careful attention because it points to a gap that is difficult to close quickly.

The Engineers Australia report provides crucial context: 78 per cent of surveyed engineers report that their AI learning is self-directed, with organisations providing minimal formal training. Yet 73 per cent recognise that AI-related skills will be essential within the next five years. This creates a readiness crisis. Without structured, employer-sponsored learning programs, the skills gap will widen. The younger cohort adopting AI faster signals both opportunity and

risk: opportunity if that capability is channelled effectively, but risk if it creates a false sense of readiness without deeper understanding.

The issue is not simply that organisations lack data scientists. It is that the broader maintenance and reliability workforce does not yet have the literacy to work effectively alongside AI tools. They need to understand what AI outputs mean, when to trust them, and when to push back. A rail infrastructure manager captured this friction:

**Understanding the bit in the middle of these AI engines is probably our biggest barrier to trusting what comes out of it. We still spend as much time checking the answer as we would have done, probably, getting to the answer in the first place.**

**Rail Infrastructure Manager**

If that remains the case, the productivity promise of AI evaporates.

This skills gap operates at multiple levels. At the leadership level, there is often insufficient understanding of what AI can and cannot do, leading to unrealistic expectations. At the engineering level, there is a need for practical AI literacy, not programming skills, but the ability

to interpret model outputs and understand their limitations. At the technician level, the challenge is workflow integration: how to incorporate AI-generated recommendations into daily routines without creating distrust or confusion.

### **Legacy Systems and Integration Challenges**

Integration with legacy systems was tied with unclear ROI as the third most cited barrier, each at 10 per cent. This is a practical, infrastructure-level problem. Most maintenance organisations run on enterprise asset management systems that were designed and implemented well before the current generation of AI tools existed. Connecting these systems to AI platforms, whether for data extraction, real-time monitoring integration, or automated work order generation, is not straightforward.

A reliability engineer from a coal operation described engaging external parties to attempt data integration for AI use cases, only to have the process stall repeatedly because of IT rules and data transfer limitations. A digital engineering manager from a major road and tunnel operator summed up the challenge:

**There's no out-of-the-box integration available, so there's a lot of work that needs to happen before data can be converted into a meaningful insight. The underlying data needs to be in a condition that it's not rubbish when the AI is providing us the output, so that people can actually trust what's been spat out at them.**

**Digital Engineering Manager,  
Road and Tunnel Operator**

This is a common pattern: organisations can identify where AI should add value but cannot bridge the gap between their existing systems and the AI tools that need access to their data.

Emerging solutions such as edge intelligence and federated learning show promise for addressing these integration and data sovereignty challenges. Rather than sending raw data to cloud systems (which raises sovereignty, latency, and privacy concerns), federated learning keeps data on-premise and distributes only the model weights across the network. This approach achieves sub-10ms latency for real-time asset decisions while respecting data residency requirements. Purpose-built edge AI processing platforms support this architecture, enabling edge-based intelligence without the bottlenecks of traditional cloud integration. This is particularly relevant for mining, critical infrastructure, and regulated industries where data sovereignty and real-time responsiveness are non-negotiable.

### **The Disconnect Between IT and Asset Management**

One barrier that deserves specific attention is the disconnect between digital or IT teams and asset management teams, selected by 8 per cent of respondents. This is not a technology issue. It is an organisational design issue. When AI initiatives are treated as IT projects, they tend to be shaped by IT priorities, IT timelines, and IT governance frameworks. These may not align with the operational realities of a maintenance organisation. MIT Sloan Management Review makes the same observation, noting that when AI predictive maintenance initiatives are treated as IT projects, they are likely to fail from the start.

The roundtable surfaced this tension directly. One participant described how their organisation had governance structures around responsible AI use driven by the corporate and technology function, while the asset management teams were trying to solve quite specific operational problems. Bridging that gap requires operational leaders to own the AI engagement alongside the IT team, not hand it over.



# Unleashing intelligence at scale for Australia's critical assets

Octave delivers context-aware, purpose-built AI to transform asset data into actionable intelligence, improving reliability, predicting risk and strengthening lifecycle performance.

# Organisational Readiness

The survey asked respondents to rate their organisation's readiness across six dimensions on a scale from 1 (Very Low) to 5 (Very High). The results provide a revealing snapshot of where organisations feel prepared and where they recognise significant gaps.

Leadership commitment received the highest score, suggesting that in many organisations, there is genuine intent at the top to pursue AI. Technical infrastructure scored reasonably, reflecting that most organisations have at least the foundational computing and network capabilities in place, even if they are not yet optimised for AI workloads.

But the two lowest scores, data quality and workforce skills, sit at the foundation of any AI initiative. You cannot deploy effective AI without adequate data, and you cannot sustain it without people who understand how to work with it. The gap between leadership ambition at 3.34 and workforce readiness at 2.69 represents the core challenge for many organisations: intent at the top is outpacing capability on the ground.

This pattern was clearly visible in the roundtable. A representative from a major mining company de-

scribed a carefully structured approach that stands out as an exception in the current landscape:

**We've done a two-year design of this digital strategy. We just started having people into the strategy this financial year. That means having KPIs and having funding top-down, endorsed by the shareholders and the board, which means we are having people cordoned off to actually intentionally work on this.**

**GM Asset Management, Mining Company**

That level of investment and organisational commitment is the exception. For most of the other participants, AI adoption was more organic, driven by individuals experimenting with available tools rather than by a coordinated enterprise strategy.

The implication is that for the majority of organisations, the path to effective AI adoption runs through investment in data foundations and workforce development, not through the purchase of more technology. The technology is already available. What is missing is the organisational infrastructure to use it well.



Leadership commitment to AI



Technical infrastructure



Culture of innovation and change



Vendor / partner ecosystem



Data quality and availability



Workforce AI skills and literacy

# Confidence, Concerns, and the Trust Question

## Confidence in AI Value Delivery

Respondents were asked how confident they are that AI will deliver measurable value for asset management in their organisation within the next three years, on a scale from 1 (Not at all confident) to 5 (Extremely confident). The average score was 3.73 out of 5.

This is a cautiously optimistic result. Nobody rated their confidence at the lowest level. Around 12 per cent scored at 2, roughly 29 per cent at 3, and the largest group at 33 per cent scored 4. The remaining 27 per cent scored 5. So the majority of respondents believe AI will deliver value, but few are highly confident about it. The tone is one of guarded expectation rather than enthusiasm or scepticism. However, the masterclass delivered an important reframing: participants recognised this as a once-in-a-career opportunity for organisations that get the strategy right. The question is not whether AI will matter, but whether each organisation will position itself to capture the value or be caught in the trap of incremental, commoditised applications.

## What Practitioners Are Worried About

The survey explored eight specific concerns related to AI in maintenance and asset management. Respondents rated each as “Not concerned,” “Somewhat concerned,” or “Very concerned.” The results reveal where the anxiety sits within the profession.

The highest-rated concern, by a clear margin, was the risk of younger engineers over-relying on AI outputs. Eighty-five per cent of respond-

ents expressed at least some concern, with 31 per cent rating it as very concerning. This is not an abstract worry. It reflects a profession where experience, judgement, and physical intuition are central to safe and effective operations, and where there is genuine anxiety about a generation gap.

This concern was echoed strongly in the MAIN-STREAM State of Asset Management report, where a participant issued a pointed warning:

**We’ve got 23 year olds passing themselves off as 20–30 year engineers. My problem with using AI, and especially in maintenance, is: do you identify using AI? I’m quite expecting in the next year or so to start seeing commercial contracts with artificial intelligence clauses that require you to have identified what’s being generated.**

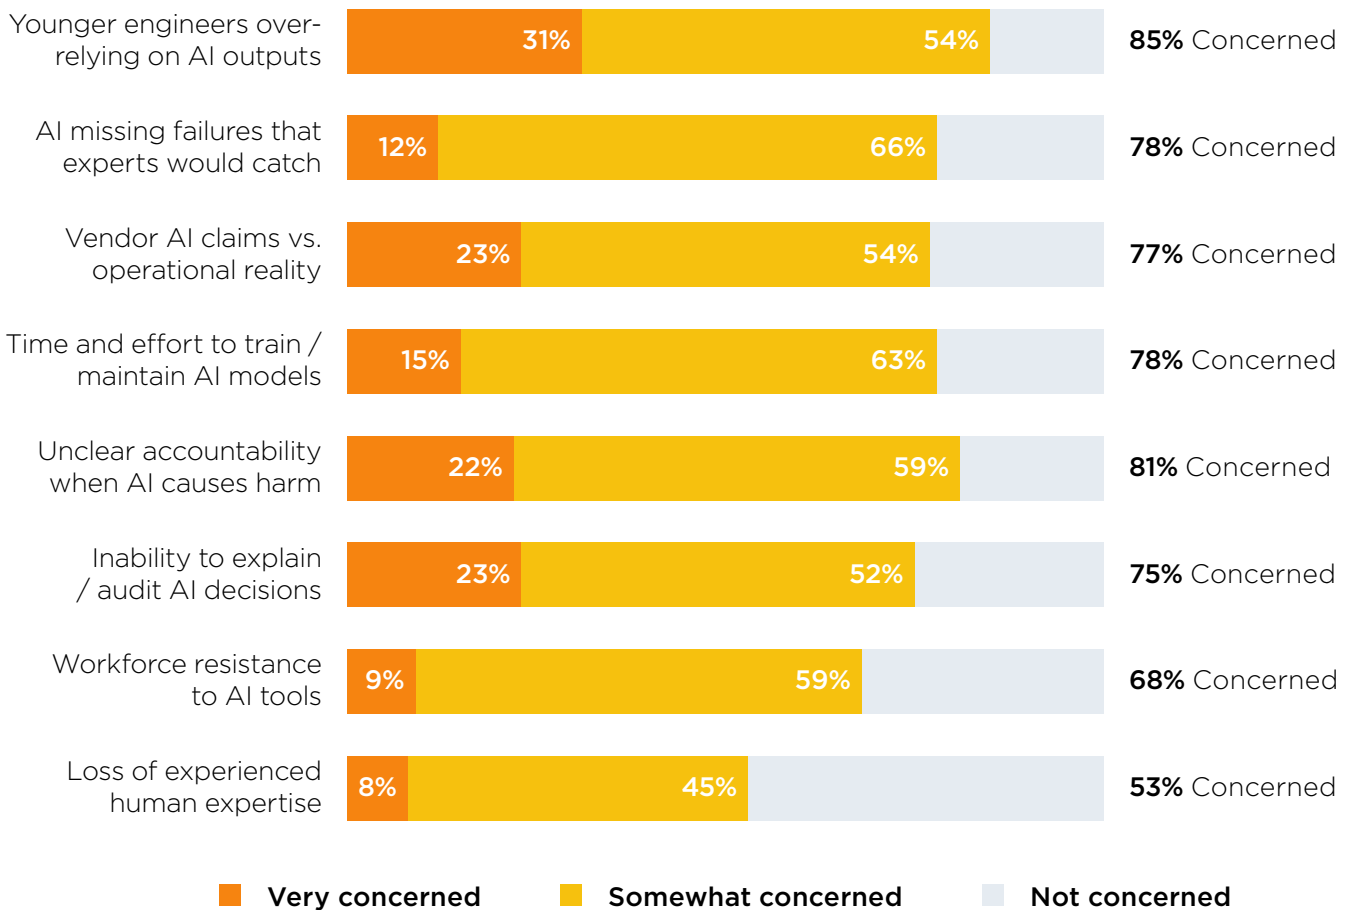
**Industry Roundtable Participant**

The fear is not that AI is bad. It is that AI, in the hands of someone who lacks the engineering judgement to evaluate its outputs, can be dangerous.

The analysis in Reliable Magazine articulates this from a technical perspective. AI does not understand physical assets. It correlates inputs and outcomes. It lacks the contextual intelligence, field experience, and failure mode understanding that only subject matter experts bring. When an AI system recommends a vibration-based intervention for equipment that does not rotate,



## What Practitioners Are Worried About



or suggests a temperature threshold no manufacturer has published, it is not lying. It is interpolating from patterns in data it does not fully understand. As the authors put it:

**AI doesn't know when it doesn't know – and that's what makes it dangerous. It can't signal uncertainty the way a human apprentice might. That's why hallucinations are so dangerous – they're confidently wrong. Without expert review, they get mistaken for insight.**

**Maintenance and Reliability  
Expert, Reliable Magazine**

The second cluster of concerns centres on accountability and transparency. More than three-quarters of respondents are worried about the inability to explain or audit AI decisions, and about unclear accountability when AI recommendations lead to harm. These are governance questions that most organisations have not yet answered, and they carry particular weight in industries operating under regulatory frameworks, such as mining, oil and gas, rail, and healthcare infrastructure.

The concern about 'AI washing' emerged from both the Engineers Australia report and the masterclass. The Engineers Australia research directly flagged AI washing as a concern, and practitioners in the masterclass noted that the term was being applied too broadly – vendors labelling any marginal automation as AI-enabled, and organisations using AI language to disguise business-as-usual. This dilutes genuine AI capability and erodes trust.

Vendor claims versus operational reality attracted concern from 77 per cent of respondents. The MAINSTREAM State of Asset Management report surfaced sharp practitioner commentary on this point:

**Every vendor promises AI will revolutionise our maintenance approach, but they never mention that their 'intelligent' system crashed our network three times in the first week because no one considered the bandwidth requirements for continuous data streaming.**

**Senior Maintenance Engineer,  
Iron Ore Operation**

Another participant described implementing machine learning for predictive maintenance only to find that technicians could not interpret the outputs, resulting in analysis paralysis rather than improved decision-making.



# Where the Opportunity Lies

Despite the barriers and concerns, the survey respondents were clear about where they see AI creating the most value over the next five years. They were asked to select up to three areas of greatest opportunity.


The top three responses, reducing maintenance costs, extending asset life, and reducing unplanned downtime, are tightly connected. They represent the traditional value proposition of better maintenance, and AI is seen as a tool to accelerate progress on all three fronts. This is consistent with the Deloitte research, which highlights how predictive approaches can deliver 25 to 30 per cent reductions in maintenance costs and 35 to 45 per cent reductions in unplanned downtime.

What is perhaps more interesting is the strong showing for knowledge capture and transfer, selected by 36 per cent of respondents. This reflects a very real anxiety in the profession about the loss of experienced tradespeople and the tribal knowledge they carry. It also connects directly to the generative AI finding: if the most widely adopted AI application is knowledge management tools, and one of the top opportunities is capturing expert knowledge, then organisations are already starting to connect those dots, even if informally.

The Reliable Magazine analysis makes a compelling case for this convergence. AI-powered platforms can link real-time asset data to instructional content, delivering just-in-time learning to technicians in the field. A short video on retorquing a flange, a visual overlay of a lubrication route, or a voice-guided walkthrough of a diagnostic flowchart. These are not theoretical possibilities. They represent a practical path to embedding expert knowledge into daily work in a way that is less dependent on which individual happens to be holding the wrench.

The same analysis also highlights a broader opportunity around standardised work. When AI is paired with structured, repeatable processes, it creates a feedback loop: consistent execution generates cleaner data, which makes the AI more accurate, which in turn makes execution more consistent. That virtuous cycle is the real prize, but it depends on having the foundational discipline in place before AI is layered on top.

## Opportunity



<span style="color: #0070C0;">■</span> Reducing maintenance costs .....	<b>54.2%</b>
<span style="color: #00BFC4;">■</span> Extending asset life / optimising capital investment .....	<b>45.8%</b>
<span style="color: #008000;">■</span> Reducing unplanned downtime .....	<b>45.0%</b>
<span style="color: #70AD47;">■</span> Capturing and transferring expert knowledge .....	<b>35.8%</b>
<span style="color: #FFD700;">■</span> Faster root cause analysis .....	<b>35.8%</b>
<span style="color: #FF8C00;">■</span> Better workforce planning .....	<b>30.8%</b>
<span style="color: #D9534F;">■</span> Improving safety outcomes .....	<b>26.7%</b>
<span style="color: #800080;">■</span> Improving regulatory compliance .....	<b>13.3%</b>

# Outcomes to Date

For organisations that have implemented AI in some form, the survey asked about the primary outcome. The responses provide a reality check.

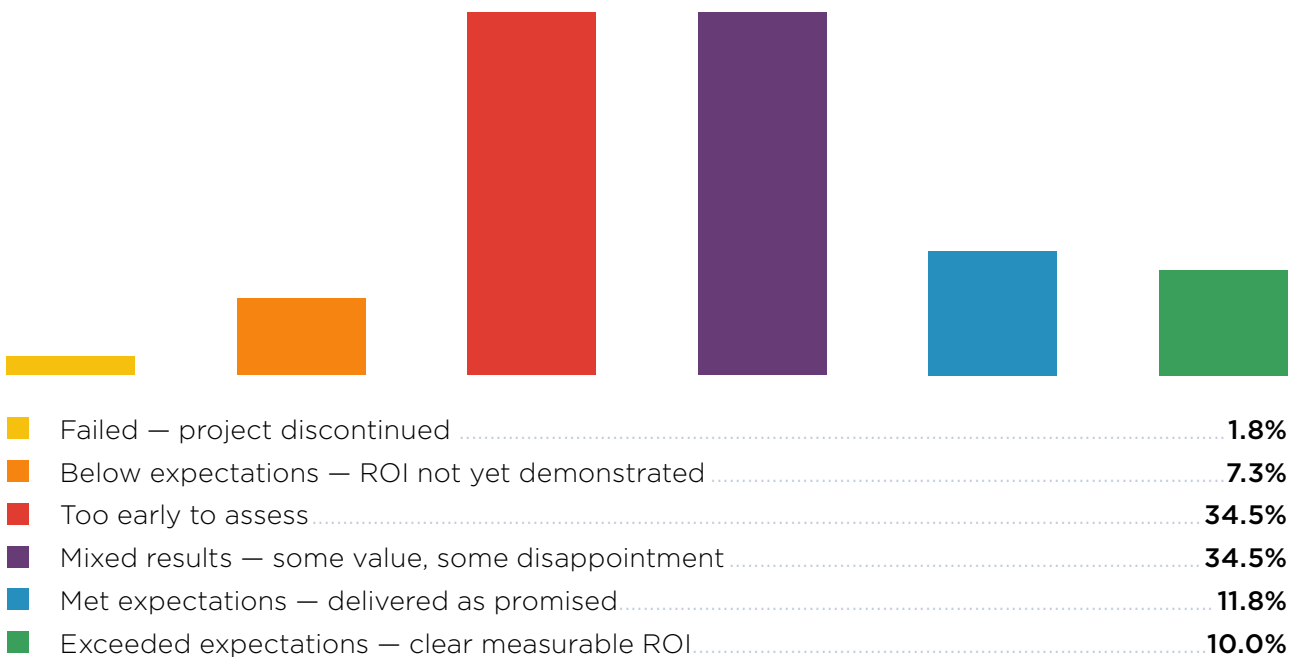
The two largest groups, each at 34.5 per cent, reported either that it was too early to assess or that results were mixed. Combined, these two categories account for nearly 70 per cent of respondents. Only about one in five reported outcomes that met or exceeded expectations. This is not a story of failure, but it is not a story of runaway success either. It is a story of organisations still working out how to make AI deliver on its promise.

The ‘mixed results’ category is the most instructive. It suggests that AI is delivering some value in specific pockets but is not yet producing the kind of consistent, scalable benefit that would justify broader rollout.

This is exactly the pattern described during the roundtable, where participants consistently reported localised wins, often driven by individual initiative, that had not yet translated into enterprise-wide impact.

The 10 per cent who reported exceeded expectations and clear measurable ROI deserve further study. Understanding what those organisations did differently, in terms of data preparation, use case selection, change management, and skills development, would provide valuable lessons for the rest of the sector. Early indications from both the roundtable and supporting literature suggest that success correlates with starting small and targeted, investing in data quality before deploying models, keeping humans in the decision loop, and having genuine operational leadership rather than treating AI as an IT-driven initiative.

## Outcome





# The Human-AI Collaboration Imperative

If there is a single theme that runs through the survey results, the roundtable discussion, and every piece of supporting literature reviewed for this report, it is this: AI in asset management works best when it augments human expertise rather than attempting to replace it.

This is not a soft, feel-good message designed to reassure nervous tradespeople. It is a hard-edged operational reality. AI does not understand failure mechanisms. It does not know what a bearing sounds like when it is beginning to pit. It cannot feel that a motor is running warmer than it should. It does not understand why a torque sequence matters for flange integrity, or why a gearbox needs to be warmed before oil is sampled. It correlates patterns in data. That is valuable, but it is not the same as understanding.

The case study cited in Reliable Magazine makes this vivid. At a mining site operating haul trucks, a predictive AI platform initially flagged numerous false positives from normal operational variations like steep inclines and sharp turns. It was only when domain experts reframed the inputs, focusing the model on steady-state conditions and excluding known transient events, that the system began delivering accurate results. Weeks later, it detected a subtle drop in oil pressure that signalled a developing pump failure, saving significant time and cost. The point: the AI detected the deviation, but the expert diagnosed it.

This pattern, where AI identifies the signal and the human interprets the meaning, is the model that works. It is what the Deloitte research calls

the collaborative approach: augmenting maintenance workflows with real-time automated monitoring, aided by human review and timely action on the ground. It is not total automation. It is intelligent partnership.

The masterclass introduced an important evolution in thinking: from ‘human in the loop’ (where humans review every AI output) to ‘human on the loop’ (where humans are positioned to intervene at critical decision points, but the system operates autonomously in routine conditions). This represents a maturation of trust and capability, but it requires even more rigorous design of which decisions truly require automation versus human judgment.

Equally important is the recognition from the masterclass that field worker disengagement has been a byproduct of heavy governance and centralised control. When tradespeople are treated as data-entry robots rather than problem-solvers, they ‘turn off their brain.’ AI offers an opportunity to reverse this through better capture of unstructured field insights (voice notes, photos, observations) and site-localised recommendations. This positions AI not as a surveillance mechanism but as a tool that amplifies field expertise.

For this to work in practice, organisations need to think carefully about how AI fits into their operational workflows. As a healthcare data and analytics leader observed during the roundtable:

**Start with the end in mind. What is it that you’re trying to achieve, measure, do? It’s that thinking that needs all the**



**work, and then all the business rules off the back of it. The technology's always easy at the other end, but the hard thing and the barrier is doing the thinking.**

**Asset Data and Analytics Leader, Healthcare**

AI recommendations need to arrive in a format that technicians can understand and act on. The reasoning behind recommendations needs to be at least partially transparent. And there need to be clear escalation paths for situations where human judgement overrides an AI suggestion, so that the system learns from those corrections over time.

The roundtable discussion surfaced a useful framing from the roundtable facilitator, who described three levels of AI deployment. Level one is general productivity tools like Copilot. Level two is applying AI agents to specific

existing processes to make them faster. Level three, and the most valuable, is rethinking the entire process design with AI capabilities in mind. As he put it:

**It's not a technology problem so much as a where-to-deploy problem. What's the cost and return, and where do we actually deploy these things? That's often why many of them don't go beyond pilot – we prove the pilot, but then can it really scale from a cost perspective.**

**Industry Roundtable Facilitator**

Most organisations are operating at level one, some are experimenting at level two, and almost none have reached level three. Getting there will require not just technology investment but a fundamental shift in how maintenance processes are conceived and managed.

# Emerging Architectures: Edge Intelligence and Federated Learning

As organisations across Australia and New Zealand grapple with AI adoption, a critical infrastructure challenge has emerged: how to deploy AI capabilities while respecting data sovereignty, managing latency, and maintaining privacy.

Traditional cloud-centric AI architectures, while powerful, introduce latency that can be prohibitive for real-time asset decisions and raise sovereignty concerns in regulated industries. Edge intelligence paired with federated learning offers a compelling architectural alternative that is beginning to gain traction.

Edge intelligence means deploying AI models at the source of data — on-site, at the edge of the network — rather than sending all data to centralised cloud systems. This dramatically reduces latency, enabling real-time decision-making critical for asset management. Federated learning takes this further: instead of centralising raw data, the approach distributes learning across edge nodes. Only the model weights are shared and aggregated back to a central repository, not the underlying data. This preserves data privacy and sovereignty while still enabling organisations to benefit from collective learning.

A leading edge intelligence practitioner in the ANZ region articulates the strategic importance. The primary barriers he identifies are four-fold: isolated knowledge silos within organisations, data silos across different systems, cloud latency that makes real-time asset decisions infeasible, and the privacy and data sovereignty concerns that are non-negotiable in mining, critical infrastructure, and healthcare contexts. Federated learning addresses all four simultaneously.

The practical implementation is becoming tangible. Purpose-built edge AI processing platforms exemplify this architecture, enabling sub-10ms latency for asset decisions while keeping sensitive data on-premise. For a mining operation managing haul trucks, or a power station optimising turbine performance, or a rail network detecting rail defects in real-time, this capability is transformative. A thermoelectric plant managing predictive maintenance for critical equipment cannot afford the 500ms round-trip latency of cloud-based decisions. A mining site managing autonomous fleet operations requires decisioning that is instantaneous and fully within their own control.

The federated learning model also has governance advantages. Each organisation retains full control of its data. Privacy-preserving machine learning is not just a technical capability; it is a governance principle that aligns with Australian and New Zealand regulatory expectations, particularly in sectors like healthcare, mining (where Indigenous land agreements may restrict data movement), and critical infrastructure (rail, electricity).

However, adoption of edge and federated architectures requires investment in edge computing infrastructure that many organisations have not yet made. It also requires a shift in how data and models are governed. Instead of a centralised data warehouse and a single AI platform, organisations need to think about distributed model governance, federated learning pipelines, and edge-cloud orchestration. This represents a maturation of AI capability that most organisations in the region have not yet reached. But for those in regulated industries, for those managing sensitive data, or for those requiring real-time decisioning, this architecture may prove not optional but essential.



# Enterprise Asset Management

AI-Powered all-in-one asset management with maintenance, document management, and lifecycle visibility in one connected platform.



## **Meridian**

EDMS for Highly Regulated Environments

## **RedEye**

Cloud Engineering Document Management

## **Maintenance Connection**

Computerized Maintenance Management Software

## **Accruent Field**

Mobile-first Scheduling and Dispatch

## **Making a Difference Across Industries**

Accruent EAM enhances asset performance, compliance, and uptime across Utilities, Manufacturing and Energy industries.

Visit [accruent.com](https://accruent.com) to learn more

# The Generational Divide and the Self-Taught Workforce

One of the most striking findings to emerge from the Engineers Australia AI-in-Practice Report concerns the generational asymmetry in AI adoption.

Across the surveyed engineering workforce, 43 per cent of practising engineers currently use generative AI in their work. Among engineering students and recent graduates, the figure jumps to 74 per cent. This 31-percentage-point gap signals a workforce in transition, with younger cohorts approaching AI as native capability while many mid-career and senior engineers are still cautiously evaluating.

The asymmetry matters because it challenges traditional hierarchies of expertise. In many maintenance organisations, the most experienced practitioners have spent decades developing intuition and domain knowledge that, properly captured, is invaluable. But they are also the cohort least likely to be regular AI users. Younger engineers, comfortable with AI tools, lack the contextual judgment to know when an AI output is plausible versus spurious. This creates a tension that organisations have not yet solved.

Equally significant is the finding on learning and development: 78 per cent of surveyed engineers report that their AI learning has been self-directed, while only 22 per cent say their organisation has provided formal training. Yet 73 per cent agree that AI-related skills will be essential within the next five years. This disconnect is a structural problem. Self-directed learning creates inconsistent capability and leaves junior staff without the scaffolding formal training provides. It also deepens the generational divide: younger engineers, digitally native and comfortable learning online, are more likely to self-educate successfully, while older practitioners may opt out.

The Engineers Australia report frames this trend with an analogy that proves instructive: 77 per cent of engineers agree that AI will become as ubiquitous as calculators once were — a tool so embedded in everyday work that its use is unremarkable. If that parallel holds, then the current moment represents a window of opportunity to ensure that the transition is managed deliberately, not left to market forces. Calculators did not replace engineers; they freed engineers from calculation drudgery. But the transition only worked because schools and organisations made a deliberate choice to teach calculator use and to redesign work around that capability.

**For asset management organisations in the ANZ region, the lesson is clear. The generational divide will not self-correct. It requires intentional investment: structured AI literacy programs tailored to different roles, intentional pairing of experienced practitioners with AI tools, and explicit recognition that generational diversity in AI capability is an asset to be managed, not a problem to be ignored.** The alternative is to abdicate responsibility for capability development to individuals and to risk a workforce where younger engineers over-rely on AI and older engineers are marginalised.

The culture of learning and openness matters enormously. The Engineers Australia report emphasises that a ‘culture of openness’ predicts faster AI adoption. Organisations that explicitly encourage experimentation, that treat failures as learning opportunities, and that invest in peer-to-peer knowledge sharing (not just formal training) see higher rates of AI adoption and more sustainable implementation. It is foundational to making AI capability stick.



# From Efficiency to Reinvention: The Case for Dynamic Asset Strategy

The masterclass revealed a critical insight that fundamentally reframes how organisations should think about AI value. Most organisations currently apply AI to process bottlenecks — planning, scheduling, turnarounds — seeking marginal efficiency gains of 5 to 10 per cent. This is the wrong target. The real value lies not in optimising existing static strategies but in reinventing asset strategy itself.

**No amount of optimisation of a static plan will help if the fundamental assumption is wrong. Most maintenance strategies operate under the same constraint. They assume a stable, predictable operational environment and optimise within those bounds. But in reality, asset utilisation, operating conditions, and production priorities shift constantly.**

The distinction matters enormously. Optimising an existing process delivers perhaps 10 per cent improvement and remains fundamentally constrained by the process's original design. Reinventing the process — making it dynamic, responsive, and tied to actual operating goals rather than budgeted resources — can unlock 2 to 3 times the value. But it requires rethinking asset management from the ground up.

The barrier to this reinvention is structural. In most organisations, the budget drives the strategy. Finance determines how much can be spent on maintenance, and asset management teams fit their plans around that constraint. AI applied to this static structure can help, but it is optimising the wrong thing. The real opportunity would be to reverse that logic: let operating

goals (availability targets, production targets, safety objectives) drive the asset strategy, and then determine the resources required. This requires breaking out of the siloed governance structures that separate finance, operations, and asset management.

The masterclass highlighted the concept of a 'unified data model' — a single, coherent picture of assets, conditions, operating context, and business objectives that connects strategy-setting all the way down to daily work execution. Current organisations typically have fragmented models: the enterprise asset management system, the CMMS, the SCADA system, the finance system, the production planning system. None of them talk to each other. Building a unified model is technically challenging but more fundamentally an organisational redesign problem. It requires breaking down data silos, establishing shared ownership, and accepting that real-time visibility into asset state and operating context is not optional but essential.

Multi-agent orchestration emerged as a critical concept in this reinvention narrative. Rather than deploying single-purpose AI models (one for predictive maintenance, one for planning, one for scheduling), orchestrated agents work together on complex, non-deterministic problems. An example from the masterclass: improving asset availability from 83 per cent to 85 per cent is not a deterministic problem with a single right answer. It requires balancing multiple competing objectives — uptime, cost, risk, workforce scheduling — and adapting dynamically as conditions change. A single AI model cannot



solve this. But multiple agents, each responsible for a different dimension and working within a shared data model, can iterate toward an optimal solution. This represents a fundamentally more sophisticated deployment of AI than most organisations have attempted.

The caution sounded during the masterclass is worth underscoring: ‘Gotta be careful we don’t cement in place anything that’s constraining.’ The history of maintenance is littered with

examples of organisations optimising themselves into corners — legacy systems that were state-of-the-art when implemented but became straightjackets. Applied to AI, the risk is that organisations will build architectures, processes, and governance around today’s AI capabilities and find themselves constrained by those choices when the technology evolves. Better to build flexibility in from the start, to assume that the AI landscape will change, and to treat AI platforms and governance as evolving rather than fixed.

# Practical Recommendations

Based on the combined evidence from the survey, roundtable, masterclass, and supporting literature, the following recommendations are offered for maintenance and asset management leaders navigating AI adoption.

## 1. Invest in data foundations

### before investing in AI models

Every data source reviewed for this report reaches the same conclusion: data quality is the single biggest determinant of AI success or failure. Organisations should audit their data environments honestly, invest in master data standards aligned to frameworks such as ISO 14224, enforce data discipline at the point of entry, and treat data governance as an ongoing operational process, not a one-off project. The principle bears repeating: 'AI won't fix broken fundamentals.' Start with reliability program discipline, then standardise, then add AI.

## 2. Treat AI like a reliability program, not a science project

This recommendation from the CEO of a leading Maintenance AI platform provider re-frames the implementation approach. Rather than treating AI deployment as an experiment that will eventually mature into something valuable, organisations should apply the rigour and discipline of a reliability program from the outset. This means clear governance, defined metrics, regular audits, and continuous improvement cycles. It means starting narrow with well-defined use cases, then systematically broadening scope only as capability matures. Trust beats sophistication.

## 3. Start with targeted, high-value use cases

Organisations that have achieved positive results consistently report starting with

narrowly defined use cases where the data is good, the business impact is measurable, and the risk of failure is manageable. Visual inspection, image recognition for defect detection, and generative AI for knowledge management are practical starting points that can build confidence without requiring perfect data infrastructure.

## 4. Build AI literacy across the workforce, with structured programs and strong organisational commitment

The skills gap identified in this research is not going to close on its own, and self-directed learning is not a sufficient strategy. Organisations need to invest in practical AI training that focuses on understanding outputs, recognising limitations, and working effectively alongside AI tools. This is not about turning maintenance engineers into data scientists. It is about ensuring they can evaluate an AI recommendation with the same rigour they would apply to any other piece of engineering advice. Programs should be role-specific, should include pairing of experienced practitioners with newer tools, and should be embedded in career development and KPI frameworks.

## 5. Preserve and embed expert knowledge while you still can

The convergence between generative AI adoption and concern about knowledge loss is an opportunity waiting to be seized. Organ-



isations should be actively using AI tools to capture, structure, and distribute the expertise of their most experienced people, embedding that knowledge into playbooks, work instructions, decision-support tools, and training materials before it walks out the door. This is not about replacing experts; it is about amplifying their impact across the entire workforce.

#### **6. Design for human-AI collaboration from the outset**

AI should be positioned as a second opinion or a diagnostic assistant, not as an autonomous decision-maker. Experienced operators should validate AI recommendations, and their corrections should feed back into the system to improve accuracy over time. This builds trust, improves outcomes, and maintains the engineering judgement that no algorithm can replicate. Evolve from 'human in the loop' to 'human on the loop' only as confidence and capability mature.

#### **7. Demand transparency from vendors and be wary of agent washing**

The gap between vendor claims and operational reality is a consistent frustration across the sector. Organisations should insist on clear evidence of outcomes achieved in comparable environments, understand the data and infrastructure prerequisites before commit-

ting, and build contractual protections around performance expectations. Be particularly sceptical of marketing language that labels routine automation as 'intelligent agents.' The distinction between real agent orchestration and process optimisation matters.

#### **8. Align AI governance with operational reality and resist the IT takeover**

AI governance frameworks should be developed with genuine input from operational teams, not imposed solely by IT or corporate functions. Accountability frameworks for AI-informed decisions need to be clear before those decisions are delegated, not worked out after something goes wrong. Most importantly, operational leaders must own AI adoption. If it is treated as an IT project, it will fail in asset management contexts.

#### **9. Avoid cementing constraints: build flexibility into architecture and governance**

Design systems, processes, and governance with the assumption that they will evolve. The AI landscape is changing rapidly, and organisations that lock themselves into rigid architectures or governance models today will find themselves constrained tomorrow. Build for flexibility, assume learning will be continuous, and treat capability as a journey rather than a destination.

# Conclusion

Artificial intelligence is not going to transform maintenance and asset management overnight. The evidence from this research is clear that the technology has genuine potential, but the organisational conditions for that potential to be realised are not yet in place for most organisations in Australia and New Zealand.

The practitioners surveyed and interviewed for this report recognise this as a once-in-a-career opportunity. Not because AI is clever, but because organisations that get the strategy right — that shift from optimising processes to reinventing strategy, that build data foundations, that invest in their people, and that treat AI as a reliability program rather than a science project — will unlock value that their peers cannot match. The question is not whether AI matters. It is whether each organisation will position itself to capture the value or be caught in the trap of incremental, commoditised applications.

The foundations matter. Clean data, integrated systems, skilled people, clear governance, and a culture that supports experimentation without punishing honest failure. These are not glamorous topics. They do not make for exciting conference keynotes. But they are the difference between an AI pilot that generates a promising dashboard and an AI capability that genuinely changes how an organisation manages its assets. As the roundtable facilitator observed:

**Tradespeople are one of the most underutilised resources we have in asset management organisations right now. We've kind of stifled innovation and creativity and continuous improvement in the field. We need to unlock that and enable those people to be part of the solution.**

**Industry Roundtable Facilitator**

The practitioners surveyed and interviewed for this report are neither fearful nor naïve. They see the opportunity. They understand the limitations. And they are asking the right questions: Where will AI add real value for us? Do we have the data to support it? Do our people know how to work with it? Can we trust what it tells us? Will this reinvent our approach or optimise our constraints?

Those are the right questions. Answering them honestly, and investing accordingly, is the path forward. The organisations that do will lead their sectors. The organisations that don't will be followers. In an industry where leadership depends on the quality of asset decisions, that distinction will compound over time.



# References

---

- Deloitte AI Institute. Using AI-Enabled Predictive Maintenance to Help Maximise Asset Value. 2023.
- Engineers Australia. AI-in-Practice Report: Skills, Adoption, and Generational Perspectives in Engineering. January 2025, n=2,197.
- CMMS Industry Leadership. Treating AI as Reliability Program: Trust Over Sophistication. Industry discussions, 2026.
- IBM. The Evolution of Maintenance — From Reactive to Predictive. 2025.
- MAINSTREAM AI Roundtable Transcript, Think Tank Series, March 2026.
- MAINSTREAM AI Survey 2026 (n=715). Survey conducted February–March 2026.
- MAINSTREAM Masterclass Series. Industry Roundtable Facilitator, on AI Adoption Strategy in Asset Management. March 2026.
- MAINSTREAM WA Report 2026 — AI Integration and Adoption Barriers extract.
- MIT Sloan Management Review. Where Traditional Maintenance Goes Off the Rails. 2025.
- Edge Intelligence and Federated Learning for Data-Sovereign AI in ANZ. Industry roundtable discussion, March 2026.
- Reliable Magazine. The Smart Way to Use AI in Maintenance and Reliability? With Experts. May 2025.
-





## Industrial Asset Management AI Report

### Contact

Ian Collier

Program Director, MAINSTREAM  
ian@mainstreamcommunity.com

Chris Carroll

CEO, MAINSTREAM  
chris@mainstreamcommunity.com

Lisa Irving

MAINSTREAM Community Director  
lisa@mainstreamcommunity.com  
+61 411 883 266

[www.mainstreamcommunity.com](http://www.mainstreamcommunity.com)