

10 AI Best Practices Every EAM Team Should Get Right Before Chasing the Hype

Here are 10 AI best practices worth getting right, based on how EAM and Maximo teams are applying AI today, not on how marketing decks describe it.

A quick note: Throughout this article, AI refers broadly to predictive AI, generative AI assistants, AI copilots, and emerging agentic AI capabilities now appearing across modern EAM platforms.

1. Fix Your Data Before You Touch a Model

Whether you're using IBM Maximo, SAP EAM, HxGN EAM, Oracle, or another enterprise asset management platform, the same rule applies: every AI capability, from predictive maintenance to natural-language work order search, depends on clean asset hierarchies, consistent failure codes, and complete work order histories. If your PM data has been inconsistent for a decade, an AI model trained on it will confidently produce inconsistent recommendations. Data quality work is not a prerequisite step you can skip to get to the interesting part. It is the interesting part.

2. Start With a Narrow, Measurable Use Case

Organizations that succeed with AI in EAM tend to start small: reducing false-positive alerts for one asset class, speeding up parts lookup for technicians, or summarizing long work order histories before a shutdown. Organizations that struggle often start with "let's use AI to transform maintenance" and never land on anything concrete. Pick one problem, define what success looks like in numbers, and prove it before expanding.

3. Keep a Human in the Loop, Especially for Safety-Critical Decisions

AI-generated recommendations, whether that is a suggested work order priority, a failure prediction, or a parts substitution, require a qualified person to review them before they drive action on critical assets. This is not a temporary training-wheels phase. In regulated and safety-critical environments, human review of AI output is a permanent part of the workflow, not a step you graduate out of.

4. Define Agent Autonomy Boundaries Before You Turn One Loose

Agentic AI is distinct from predictive AI (which forecasts), prescriptive AI (which recommends), and generative AI (which drafts or summarizes) because it does not merely forecast or recommend; it acts. An agent can create a work order, reorder a part, adjust a PM schedule, or escalate an alert on its own, chaining several steps together without a person clicking approve at each one. That changes the governance question. It is no longer "should someone review this recommendation?" but "how much can this agent do before a human sees anything at all?"

Before deploying an agent in a Maximo or EAM environment, get clear answers to a short list of questions: What actions can it take without approval, and which always require sign-off? Is there a complete audit trail of every action it takes, not just the ones that go wrong? What is the rollback path

if it acts on bad or stale data? And critically, does it escalate when it is uncertain, or does it guess and move forward anyway? An agent that cannot tell you when it does not know something is not ready for autonomy in a safety-critical environment.

5. Treat Technician Trust as a Design Requirement, Not an Afterthought

A model can be statistically accurate and still get ignored on the floor if technicians do not trust it. Trust is built by showing the reasoning behind a recommendation, not just the output, and by being transparent when the model is wrong. Involve technicians and planners in evaluating AI tools before rollout, not just in training after the decision has already been made.

6. Audit Your AI Vendor's Training Data and Model Governance

Ask where the model was trained, whether your organization's data is used to train models that other customers benefit from, and what happens to your data if you switch platforms. These are reasonable procurement questions in 2026, not paranoid ones. A vendor who cannot answer clearly is telling you something.

7. Set Realistic Expectations for Accuracy, Especially Early On

Early-stage predictive models in EAM are often right more than they are wrong, but they are not perfect, and false positives are common while a model is still learning your environment. Set expectations with stakeholders up front so a few early misses do not undermine confidence in a tool that would improve with more data and time.

8. Map AI Initiatives to Your Existing Reliability Strategy

AI is not a reliability strategy on its own. It is a tool that supports one. If your organization already has a clear RCM or reliability-centered approach, AI should fill gaps in that strategy, such as surfacing patterns humans would take too long to find manually. If you do not yet have a reliability strategy, building one should come before layering AI on top of the gaps.

9. Plan for Change Management Like It Is Its Own Project

New AI features change how planners triage work, how technicians receive instructions, and how supervisors review performance. Each of those role changes requires its own training, communication, and feedback loop. This matters even more with agentic AI, since an agent acting on someone's behalf changes not only their workflow but also their sense of control over it. Treating an AI rollout as a software update rather than a workflow change is one of the most common reasons adoption stalls after a promising pilot.

10. Measure What Actually Matters, Not What Is Easy to Report

Model accuracy percentages look good in a vendor case study. What matters to your organization is whether unplanned downtime dropped, whether technician time spent on diagnostics went down,



and whether parts inventory carrying costs improved. Define the operational metrics that matter to your business before rollout, and hold the AI initiative to those metrics, not to the ones that were easiest to pull from a dashboard.

The Bottom Line

AI in EAM is not a single decision. It is a series of smaller decisions about data, trust, governance, and change management that compound over time. The organizations that get real value from it are not the ones with the flashiest pilot. They are the ones who did the unglamorous work first: clean data, clear use cases, human oversight, and metrics tied to actual operational outcomes.

If your team is somewhere in the middle of that process, or still deciding where to start, that is a normal place to be. The gap between the AI hype cycle and what is working on the floor right now is wide, and closing it is worth the deliberate pace.

