



AI-readiness assessment for IBM Maximo estates



Table of Contents

Why AI-readiness is a Maximo question, not just an AI question.....	3
Dimension 1 — Data quality.....	4
Dimension 2 — Criticality model.....	4
Dimension 3 — Integration.....	5
Dimension 4 — Governance.....	6
Dimension 5 — Operating model.....	6
Putting it together — a one-page readiness profile.....	7
Where Maxiron fits.....	7

Document ID	MXR-GUIDE-007
Version	1.0
Date	2026-04-24
Owner	MaxIron Ltd
Audience	Procurement, IT, asset management and operations leaders evaluating IBM Maximo and MAS
Classification	Public — share freely with credit
Canonical URL	https://maxiron.com/guides/ai-readiness-assessment-for-maximo-estates/

The conversation about AI in asset management has shifted in the last twelve months. Two years ago, the question was “should we?”. Today, the question is “are we ready?”. The honest answer for most IBM Maximo and MAS estates is “not yet, and here is what the gap actually looks like.”

This guide is the framework MaxIron uses with clients to assess that gap before any AI engagement starts. It is deliberately practitioner-grade: the questions are the ones a senior consultant would ask in a working session, with the evidence we actually look for and the failure modes we expect when the evidence is not there. There are five readiness dimensions: **data**, **criticality**, **integration**, **governance** and **operating model**. A platform that scores well on all five is ready to land AI productively. A platform with even one weak dimension can land AI but will spend most of its first programme paying for the gap rather than capturing the value.

1. Why AI-readiness is a Maximo question, not just an AI question

Almost every AI use case in asset management — anomaly detection on equipment health, predictive maintenance, in-application assistants, automated work-order classification, visual inspection — depends on three things: clean asset data, an honest criticality model, and a feedback loop that puts the AI’s recommendation back into the system of record where the work actually happens.

All three of those live in Maximo. If the asset hierarchy is wrong, the AI is detecting anomalies on the wrong assets. If the criticality model is aspirational, the AI is prioritising the wrong work. If the feedback loop is broken, the AI’s recommendations land in a dashboard nobody opens. The platform is the gating constraint; the model is rarely the gating constraint.

This is also where most failed AI programmes in asset-intensive operators go wrong. The AI vendor pitches a model. The buyer signs. Six months later the model is technically working, but the data feeding it is not trustworthy, the criticality model is twelve years old, and the maintenance planners do not yet treat the model output as authoritative. The model gets quietly dropped. The investment is written off. AI in Maximo gets a bad reputation internally for two years.

The whole point of the assessment below is to surface those gaps before contract, not after.

2. Dimension 1 – Data quality

AI is not magic. It does not infer truth from contradictory data. The data quality bar for AI is meaningfully higher than the bar for traditional reporting, because traditional reporting tolerates being approximately right at the aggregate level whereas AI is making per-asset, per-event decisions where wrong data produces wrong decisions.

The questions and the evidence:

Question	Evidence we look for	What “not yet” looks like
Is the asset hierarchy modelled correctly down to the level the AI use case needs?	A documented hierarchy aligned to ISO 14224 (or sector equivalent), with locations, parent-child consistency, and a known position for every operating asset.	“Most assets are in there somewhere” or three different hierarchies running in parallel by site.
Are failure codes used consistently and is the long tail of “Other” small?	A failure-code library mapped to ISO 14224 failure modes / mechanisms, with the share of “Other” / unclassified codes under 10% and falling.	Hundreds of locally-coined failure codes, “Other” growing year on year, no shared definitions.
Is master data (vendors, items, equipment specs) deduplicated and current?	A documented MDM process, a known duplicate rate, and a known stale-data rate (for example, items not used in 24 months).	The same vendor with seven name variants, the same valve modelled five times, no inventory of the staleness problem.
Are work orders closed out with usable data, not free-text?	A high rate of structured close-out fields (failure code, failure mode, time-on-tool, parts used) on completed WOs.	Close-out comments are the only data captured; every analytic query starts with a regex over free text.

If three or more of these are red, the right first investment is **not** AI. It is data quality. MaxIron’s [AI Smart Data](#) is the product we use here because it does the heavy lifting at scale with audit trail, but the principle is the same regardless of tooling: fix the data before paying the AI vendor.

3. Dimension 2 – Criticality model

Almost every useful AI output is “rank these things by something”. The “something” is almost always derived from criticality. If the criticality model is missing or aspirational, the ranking is meaningless.

The questions:

- **Is there a criticality score on every operating asset, or only the headline assets?** A model that only covers the top 20% leaves the AI guessing on the long tail. A model that covers everything but with the same score on most things has the same problem.
- **Does the criticality model reflect current operating reality, not the design assumption from ten years ago?** Plants and networks change. A criticality model that has not been re-baselined since the last big asset acquisition or decommissioning is probably wrong in important places.
- **Is the criticality model the one that operations actually uses to prioritise work?** A “shadow” criticality model in operations and an “official” model in Maximo means the AI is optimising against a model nobody believes.
- **Is there an explicit consequence dimension (safety / environment / production / cost), not just a probability dimension?** A consequence-blind model converges on the same recommendations as no model at all.
- **Is the model auditable?** Can an engineer ask “why is asset X scored at 8?” and get an answer in evidence? If not, the model is a black box and the AI’s outputs that depend on it will be a black box too.

The Maxiron position is that any AI engagement on a platform with a weak criticality model spends its first quarter rebuilding the criticality model. There is no shortcut. We tend to use [Maximo Health and Asset Criticality](#) as the underlying framework here.

4. Dimension 3 – Integration

AI in Maximo rarely lives entirely in Maximo. The model needs telemetry from the operational stack (historian, SCADA, IoT platform), it needs context from finance and supply chain (cost, parts availability), and its recommendations need to land back in Maximo as actionable work in flight that planners can review and accept.

The questions:

- **Is there a clean read path from the historian / SCADA / IoT into MAS Monitor or an equivalent?** This is the single most common gap. We have written about it specifically in [Connecting an OT historian to MAS Monitor without breaking the OT team](#).
- **Is the integration governed?** Authentication, retry behaviour, error volumes, observability, certificate rotation. An AI model fed by an integration that silently drops messages produces wrong outputs and the data team finds out later.
- **Does the AI’s recommendation come back into Maximo as a structured artefact?** A work request with a reason code, a flag on a work order, a notification on an asset — not a CSV dropped into a SharePoint folder.
- **Is there a clear contract for what the AI is allowed to do automatically vs what requires a human?** Auto-create work requests is usually fine. Auto-close work

orders is usually not. Where that line sits has to be explicit before the AI goes live.

If integration is weak, AI lands as a parallel system that says contradictory things to Maximo. Operations stops trusting both. The fix is structural.

5. Dimension 4 – Governance

AI brings governance questions traditional Maximo work does not. Some of them are regulatory (the EU AI Act, sector-specific safety regulators). Most of them are operating-model questions the business has not yet had to answer.

The questions:

- **Who owns the model?** A model with no owner gets stale and nobody notices. The owner is usually a senior engineer on the asset side, not in IT.
- **Who reviews the model output before it becomes work?** This is rarely “no one”, but it is often “the planner, implicitly, by accepting or rejecting the recommendation”. That is acceptable if the planner has the context to push back. It is dangerous if they do not.
- **How are model performance, drift and false-positive rates monitored over time?** A model that was 92% accurate at deployment and is 71% accurate eighteen months later is producing different work. The business needs to know.
- **What does the audit trail look like?** Particularly for safety-critical decisions: which model version, which input data, which recommendation, which human accepted or rejected it. [Maxiron Change Control](#) is one of the patterns we use here, but the principle generalises.
- **Is there a documented position on data residency and on what data may leave the platform to train or fine-tune external models?** This is increasingly an audit and procurement question.

A platform that has never had to govern any model is not “ahead” — it is unprepared. The first AI use case is the one that forces these questions to be answered. Better to answer them on a small, low-risk model than to answer them in production on a high-stakes one.

6. Dimension 5 – Operating model

The hardest readiness dimension is also the least technical. Are the people who will receive the AI’s outputs ready to act on them, and are they incentivised to act on them?

The questions:

- **Will the maintenance planner trust the AI’s recommendation enough to actually re-plan the week?** If the planner’s KPIs reward planned-work compliance and the AI recommends breaking the plan, the planner will quietly ignore the recommendation. This is the single most common failure mode.

- **Will the engineer in the field accept that an algorithm prioritised this job over the one they thought was urgent?** The change management around this is harder than the AI itself.
- **Is there a clear escalation path when the AI gets it wrong?** And it will, sometimes spectacularly.
- **Has the operating model been updated to reflect AI-augmented decisions?** Procedures, roles, training, KPIs. If not, the AI is operating in an organisation that has not yet decided how to use it.
- **Are leaders prepared to defend the AI’s recommendations, or will they retreat at the first uncomfortable result?** Executive air cover for an AI programme is non-optional.

We have written more on this in [AI-led changes to the asset operating model](#) and the related insights in the AI cluster.

7. Putting it together – a one-page readiness profile

The way to use the five dimensions is to assign each one a colour: **green** (ready), **amber** (will work but with caveats), or **red** (the AI use case will fail until this is addressed). A single one-page profile gives the senior sponsor the truth in 30 seconds:

Dimension	Status	Top blocker	Recommended first action
Data quality			
Criticality model			
Integration			
Governance			
Operating model			

The honest read on most estates we assess is: one or two greens, two or three ambers, one red. The right next step is almost always to fix the red before signing the AI contract, not after. Programmes that try to do both in parallel either run twice as long or have to throw away the model six months in.

8. Where MaxIron fits

We run this assessment as either:

- A **focused readiness review** — five working days, single estate, output is the one-page profile and a remediation roadmap. This is what most clients commission first.

- A **scope add-on inside a [Maximo Health Check](#)** — when the buyer is already commissioning a wider audit and wants AI-readiness as a dedicated dimension within it.
- A **full programme** when the foundation work (data quality, criticality, integration) needs to be done before AI lands at all. This is where [AI Smart Data](#), [Cloud Manager](#), and [Change Control](#) earn their keep — they are the products that turn the readiness findings into the platform that lets AI work.

If you are weighing up an AI investment on a Maximo or MAS estate, the cheapest version of this conversation is a free 30-minute Maxiron Health Check review. Bring two or three things you would like the AI to do; we will walk through the readiness implications honestly, and tell you whether the AI is the right next investment or whether one of the five dimensions needs work first.

[Book a 30-minute review](#) · [Read the full Maximo Health Check & Heal service page](#) · [See AI Smart Data](#) · [See the AI in Maximo webinar](#)