

Beyond the alert: operationalizing human judgment with AI

Words: Ben Torkian, PhD, Chief Technology Officer, SkySpecs and Joel Saxum, VP, Growth, Partnerships and Customer Success, SkySpecs

Wind operators have more asset data than ever, but turning it into a maintenance plan still depends on engineering judgment and knowledge scattered across multiple systems. A new approach combining connected data, established engineering methods and AI could help preserve that expertise and put it to work across an entire fleet.

A condition-monitoring alert is useful, but it is not yet a maintenance plan.

Imagine an operator with several gearboxes beginning to show signs of degradation. None has failed. Some may continue operating for years; others could become expensive events if the warning signs are misunderstood or acted on too late.

The question is not simply which gearbox looks worse today. It is which components are most likely to require intervention in the next six to eighteen months, what the consequence of waiting might be and how that work should be planned across the portfolio.

Answering that question involves far more than a condition monitoring system (CMS) screen. A drivetrain engineer may review diagnostic outputs alongside supervisory control and data acquisition (SCADA) and performance data, work orders, component



history, inspection findings, parts availability, crane logistics and the cost of similar work.

The asset manager then has to translate the technical picture into a campaign, a budget and a decision about risk.

Now imagine asking the fleet a broader question: which gearboxes should we be planning to touch in the next six to eighteen months, and why?

An AI-assisted workflow could assemble an initial portfolio view, showing the evidence behind each recommendation and grouping likely interventions by timing, geography and campaign opportunity. It might also flag a unit that should not be ranked at all: 'Turbine 82 cannot be compared with the current group. Gearbox replacement records and monitoring history do not align. Engineering review required.'

The drivetrain engineer reviews that exception, challenges several assumptions and adjusts the recommendation for two other units. The operator has not delegated a major-component decision to AI. What has changed is the amount of technical and operational context that can be assembled before the team begins planning the work.

The fleet is hypothetical, but the building blocks are not. Bringing them together requires three layers: access to the relevant systems, a defined engineering method and AI that can use both to support the people who remain accountable for the decision.

Layer one: access

The gearbox question reaches across systems that were rarely designed to work as one. Condition-monitoring platforms hold component-health signals. SCADA and performance systems provide operating context. Maintenance systems hold work orders and component history. Other systems may contain inspection results, parts and spares information, contractor availability, crane planning and historical costs.

Wind operators have spent years digitizing these areas, yet much of the information still sits in separate applications. A person has often been the integration layer. In wind, the middleware frequently holds an engineering degree.

Model Context Protocol (MCP) provides one way to make that access more consistent. It gives AI applications a common interface to external tools and data sources, while the connected systems retain control over permissions and actions. That can reduce the need for bespoke connections between every AI workflow and every condition-monitoring, maintenance, inspection or planning system. Most systems may remain read-only, while people retain approval authority.

Access, however, is only the first layer. Giving a language model access to ten systems does not automatically give it a drivetrain engineer's judgment.

MCP lays the pipes. It does not provide the engineering method.

Layer two: know-how

The portfolio ranking depends on the method used by the people who understand how component-health signals become maintenance decisions.

Call the experienced drivetrain engineer Elena. She knows that a remaining-useful-life estimate is not a decision on its own.

How confident is the diagnostic? Is the trend persistent? Does the monitoring history belong to the component currently installed? What maintenance has already been performed? What is the consequence of waiting, and what missing evidence should trigger further engineering review before a major-component campaign?

Good engineering is not simply knowing what to do. It is knowing when the available evidence is not good enough to make the call.

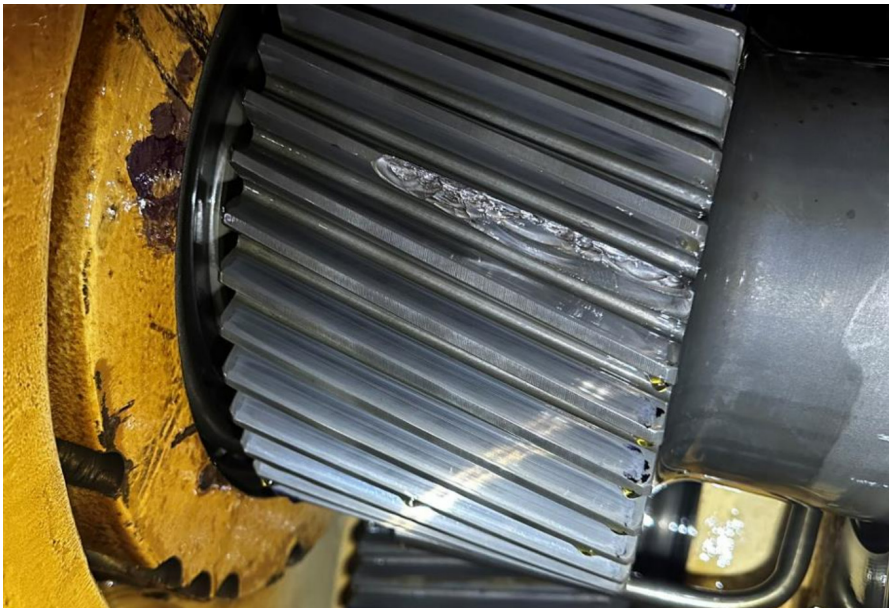
One useful way to capture that judgment for AI is as a skill. In this context, a skill is a



Ben Torkian



Joel Saxum



versioned set of instructions, criteria, references and escalation rules for a specific task. It describes the sequence to follow, the tools to use, the checks to perform and the points at which a human decision is required.

MCP provides access to the systems. The skill defines how the information should be used.

Layer three: intelligence

The third layer brings the systems and engineering methods together. For the gearbox question, AI can retrieve the relevant evidence, apply the appropriate diagnostic and planning methods and record the basis for each recommendation.

It can rank assets by likely intervention window and consequence, then consider the practical factors that determine whether a portfolio plan is executable: geography, crane timing, parts availability, previous costs and other planned work at the site.

Turbine 82 triggers an escalation because the component identity and monitoring history do not align. Rather than forcing a ranking, the workflow leaves the unit unresolved and refers it to an engineer.

That distinction matters. The most useful AI in renewable asset management will not be the system that always produces an answer. It will be the one that understands which method applies, can show the evidence it used and recognizes when the evidence is not sufficient.

Permissions stay in software. Accountability stays with people.

And now the part that changes the operating model: Elena is no longer at the company. She moved to another organization, but the method she helped formalize still supports the fleet.

From diagnosis to campaign

Once Turbine 82's records are resolved, the asset manager can move from a list of component-health concerns to the questions that determine next year's work. Which interventions fall into the same period? Which sites could share a crane campaign? What have comparable replacements cost? Where does delaying work materially change the risk?

AI can assemble and test those scenarios, but the final call remains with the engineering and asset-management team. The aim is to make the trade-offs visible rather than leave them buried across several systems and spreadsheets.

The same approach can support other parts of the fleet. After a lightning event, a team could combine strike data, sensor information and inspection history to identify which turbines justify inspection first. A blade team could examine where a defect has appeared elsewhere, how similar damage progressed and which repairs held.

AI does not suddenly acquire wind engineering expertise. But access to the right asset data, analytical tools and established methods can help engineers apply their judgment across more cases. Asset managers, in turn, gain a clearer picture of the operational and commercial choices available.

How the method improves

Every case can contribute something beyond the immediate decision. An engineer may override a ranking, a physical inspection may reveal something the monitoring data did not, or a component may perform differently than expected after an intervention.

Those outcomes can be reviewed and used to refine the method. The improvement is not necessarily in retraining the underlying model. It may come from changes to the operating method around it, which engineers can inspect, test and update.

The real value goes beyond giving engineers another way to access existing software. Knowledge that once lived in individuals, inboxes, spreadsheets and recurring meetings can become part of the organization's operating system.

That matters in an industry where experienced people retire, change employers or cannot be in every room at once.

The time is now

At SkySpecs, this is the direction we are actively building toward. We already have MCP-based capabilities in beta across blade and condition-monitoring workflows, and we are beginning to open those environments to selected operators while developing the skills, permissions and auditability around them. Broader releases will follow over the coming quarters.

The aim is not to replace the engineer, automate turbine control or create another black-box recommendation engine. It is to reduce the time spent assembling information, make expert methods reusable and give operators a clearer path from signal to decision.

Developing this approach requires input from across the industry. Operators hold the outcomes, service providers hold field evidence, engineers contribute judgment and software teams provide the infrastructure.

For operators interested in shaping how this develops, beta workflows offer an opportunity to test recommendations alongside existing processes, challenge the results and identify cases where the system should stop rather than produce an answer.

At SkySpecs, we are building the same common architecture across the broader renewable asset environment: wind blades, drivetrain CMS and performance monitoring, solar, transmission-line inspection and, over time, other asset classes such as battery energy storage systems (BESS). The analytics and engineering methods differ, but the underlying challenge is the same: connecting evidence, expert know-how and intelligence so operators can move from signal to decision with less friction.

Organizations that begin capturing those methods today will have more than another year of data next year. They will have accumulated decisions, corrections and tested judgment to draw on.

Return to the gearbox question. As the component-health record changes, the organization can continue to draw on the reasoning and engineering judgment behind earlier decisions, even when the people involved have moved on.

People move on. Their knowledge should not have to.

About the authors

Ben Torkian, PhD, is Chief Technology Officer at SkySpecs, specializing in applied AI, large language models and multi-agent systems.

Joel Saxum is VP, Growth, Partnerships and Customer Success at SkySpecs, working with renewable operators on asset-management strategy, customer outcomes and the commercial application of new technology.

skyspecs.com