

TECH & AI CASE STUDY - ENERGY / UTILITIES - 2026

IoT Predictive Maintenance That Caught Failures 3 Weeks Early

Based on a delivery engagement with Cascadia Ridge Energy

up to 3 weeks
FAILURE WARNING LEAD

-42%
UNPLANNED DOWNTIME

-27%
MAINTENANCE COSTS

94.6% to 97.8%
FLEET AVAILABILITY

A fillable practice edition: the full story, 8 actions,
6 reflections and a 90-day plan - work it with your team.

Free - no email required
emersontechnology.com

MAKE THIS PRACTICE BOOK YOURS

Work it, do not just read it

This practice book turns the Cascadia Ridge Energy case study into a working session for your team. Read the full story first - what happened, what was built and what it changed. Then tick the 8 actions you can commit to, write short answers to the 6 reflections, and finish with a 90-day plan. Every field in this PDF is fillable: open it in Adobe Acrobat Reader, Chrome, Edge or macOS Preview, type directly into the boxes and save.

YOUR NAME

ROLE

COMPANY

TEAM / DEPARTMENT

DATE

EMAIL (OPTIONAL)

MY GOAL FOR THIS PRACTICE SESSION

- 01

The story - What happened at Cascadia Ridge Energy - before, during and after the project.
- 02

The results - The numbers that changed, plus industry benchmarks from public research.
- 03

The 8-action checklist - Tick boxes as your organization completes each step. Owner and date fields included.
- 04

The 6 reflections - Fillable questions that pressure-test your own roadmap.
- 05

My 90-day plan - Three fillable phase cards that turn this case study into your roadmap.

RUN IT AS A 90-MINUTE TEAM SESSION

25 min READ Read the story together. Ask everyone to note the one number that surprised them most.	40 min PRACTICE Work the 8-action checklist and the 6 reflections. Assign an owner and a date to every action you commit to.	25 min COMMIT Fill in the 90-day plan. One goal per phase, one number that tells you it worked. Book a review date before you leave the room.
--	--	---

THE STORY

What happened, and why it worked

01 The Company

Cascadia Ridge Energy operates 96 wind turbines across four wind farms in remote terrain, selling generation under contract pricing that depends on documented fleet availability. The economics of wind maintenance are dominated by two facts the company could not escape: turbines are tall, and sites are far apart. A gearbox or bearing repair requires a crane crew mobilized to a ridge hours from the nearest service center, and every day a turbine stands still is revenue lost at the same time as repair costs accumulate. In an industry where the resource itself is free but uptime is everything, unplanned downtime is the most expensive category of cost a fleet operator carries.

The company's maintenance organization was competent, experienced and structurally blind. Technicians knew the machines well, the operations team tracked availability diligently, and the finance team could price an outage to the dollar — but nobody knew when the next failure was coming, because the maintenance model was calendar-based: inspect every 180 days and hope for the best. Regulatory reporting added pressure, since the utility had to document fleet availability to justify its contract pricing, which turned every unplanned outage into a question the company had to answer in public.

02 Where Things Stood

Before Emerson Technology engaged, the fleet ran on scheduled inspection rounds and reactive repair. Gearbox and bearing failures arrived without warning, and each unplanned outage cost roughly \$38,000 in repairs, crane mobilization and lost generation — more than \$1.1M annually across the fleet. The prior year's worst event, a catastrophic gearbox failure, took a turbine offline for 34 days, a span dominated not by the repair itself but by crane scheduling, parts logistics and the queue of consequences that follows a structural component failure in a remote location. Truck rolls to sites three hours apart wasted technician hours on healthy machines, because the inspection calendar, not the machine's condition, decided when people drove out.

The information the maintenance team needed existed only in fragments. Work orders lived in the CMMS, scada data lived in the turbines' own historian, and the Oil quality reports arrived as PDFs from a lab. Nothing joined a vibration reading to a bearing replacement to the run hours that preceded it, so the fleet's six years of failure history — an asset worth real money to anyone who could learn from it — sat unusable in disconnected systems. The 34-day outage was not an anomaly; it was the predictable output of a system that treats all machines as identical and all failures as surprises.

The financial trajectory made the status quo untenable. Availability targets were tightening as contract renewals approached, the failure-prone population of gearboxes and bearings was aging in place, and every escalation of maintenance spending without a corresponding availability gain would show up directly in the economics the regulators reviewed.

03 The Challenge

The challenge was to convert six years of scattered failure history and a remote, instrumented-just-enough fleet into predictions reliable enough to reorganize maintenance operations around — because a prediction nobody schedules against is a novelty, not a capability. The bar was high for good reason: a false alarm that mobilizes a crane crew to a healthy turbine destroys credibility faster than ten correct warnings build it, and a missed warning on a gearbox costs \$38,000 plus whatever the schedule contamination costs downstream. The system had to earn the trust of maintenance planners who had spent careers developing judgment and had every reason to protect it.

The physics and the geography both constrained the solution. Bearing wear, imbalance, overheating and oil degradation each announce themselves in different measurements, so sensor selection had to follow the OEM's failure-mode documentation rather than a generic kit. Telemetry had to traverse remote terrain where connectivity is not guaranteed. And the economics only work at the system level: individual predictions matter less than the ability to bundle repairs into planned windows that share crane and travel costs, which meant the output of the models had to feed a planning workflow, not just a dashboard.

04 Discovery and Diagnosis

Emerson Technology began with a three-week discovery sprint built around the fleet's failure history. The team reconstructed six years of work orders and failure logs from the CMMS and turbine historians, standardizing component naming and fault codes that had drifted across sites and shifts — an unglamorous exercise that turned out to be the foundation of everything that followed, because failure models are only ever as good as their labeled history. The analysis ranked components by failure cost and downtime impact and found the expected concentration: gearbox and main bearing failures accounted for the largest share of outage dollars, which is where instrumentation would start. McKinsey Global Institute analysis estimates that IoT-based predictive maintenance can reduce equipment downtime by up to 50% and cut maintenance costs by 10-40%, and US Department of Energy guidance reports that a functioning predictive maintenance program can yield savings of 30-40% over reactive approaches — benchmarks the team used to sanity-check the business case rather than to promise from.

The diagnosis also quantified the operational drag that made the case unambiguous. Technician hours consumed by driving to healthy machines, the crane mobilization premium embedded in every unplanned repair, and the schedule contamination that followed each surprise outage together formed a cost structure that no amount of inspection diligence could fix, because it was designed into the calendar-based model. Deloitte's analysis, reporting that predictive maintenance programs typically reduce breakdowns by around 70% and lower overall maintenance costs by 25%, matched the shape of the opportunity the failure data was showing. The discovery concluded with a deliberately scoped phase one: sensor the worst offenders first, prove the warning lead time, and wire predictions into the CMMS before expanding coverage.

05 The Solution

Emerson Technology instrumented the fleet and built the decision pipeline that turns telemetry into scheduled work. The architecture has three layers.

Sensing the fleet

Turbines received vibration, acoustic, temperature and oil-particle sensors chosen against documented failure modes rather than installed as generic kits — bearing wear, imbalance, overheating and oil degradation each demand different measurements. Telemetry streams through MQTT into an AWS IoT Core gateway and onward to a time-series store, with buffering for the sites where backhaul connectivity is intermittent. The sensor layer was

designed for the crane economics from the start: the goal is not data collection but early detection of the specific degradations whose repair windows can be moved from emergency to planned, and every sensing decision was justified against that test.

The models

Feature pipelines in Python compute spectral signatures — bearing fault frequencies, gear-mesh harmonics, thermal drift — that feed gradient-boosted and autoencoder models trained on the six years of standardized work orders and failure logs. Gradient-boosted models capture the labeled patterns the history documents; autoencoders catch the anomalies that have no clean label yet, learning what "normal" sounds like for each turbine and flagging departures from it. Alerting is tuned by severity, because a few high-precision warnings beat constant noise: every false alarm spends credibility the program will need for the big catch. A digital-twin comparison view lets reliability engineers place a suspect turbine alongside its fleet twins — units of the same model in similar terrain — which separates true anomalies from machines that simply run hot or sit in harsher wind, and contextual baselines like these cut false positives dramatically compared with fixed thresholds.

From alert to work order

When a model's confidence crosses threshold, Airflow triggers a work order carrying the flagged component, severity and recommended intervention window, pushed directly into the CMMS — a prediction without a planning workflow is just a novelty, so the integration was treated as the product's core rather than as a connector. Grafana dashboards give reliability engineers fleet-wide health scores, and alerts flow to mobile devices with historical sensor context attached, so field technicians arrive with the right parts the first time instead of diagnosing from scratch on a ridge two hours from the depot. Finally, the platform translates everything upward: availability, generation and dollar impact reported to executives in the language the regulatory filings already use, because programs funded on avoided-outage dollars survive budget cycles that dashboards full of model-accuracy metrics do not.

06 The First 90 Days

Weeks one and two standardized the historical record: six years of work orders cleaned, component names unified, failure causes coded, and the worst-offender population identified across the four wind farms. Weeks three through eight instrumented the first tranche of turbines — the highest-cost machines by failure history — and began streaming telemetry, with the first month spent validating that sensor placements actually observed the failure modes they were chosen for and that backhaul buffering survived the sites' connectivity gaps. Weeks nine and ten trained and back-tested the models against the standardized history, deliberately scoring them on cases the training had never seen.

Weeks eleven and twelve wired the operational loop. Alert thresholds were set conservatively by severity, the Airflow-to-CMMS trigger was tested end to end with the maintenance planning team, and the first flagged component — a main bearing showing the earliest spectral signature the team had trained on — was scheduled into a planned window rather than an emergency response. The maintenance planners set one rule for the transition: every alert would be logged with its outcome, confirmed or dismissed, and reviewed monthly. That review loop did two things — it tuned thresholds against reality instead of theory, and it gave the planners visible evidence of precision before they were asked to reorganize crane schedules around model output.

The quarter closed with the bundling economics already visible: two flagged components on the same farm were combined into a single planned crane visit, sharing mobilization costs that two emergency calls would have doubled. Nothing about that moment was technologically dramatic, and everything about it was the point — the prediction only became money when planning could act on it. By day 90 the program had its first confirmed catch, its first shared-cost repair window, and a precision record the planners were willing to schedule against.

07 The Results

The first year delivered results the fleet could feel and finance could audit. The system flagged six component failures up to three weeks before breakdown, letting the team bundle repairs into planned windows instead of emergency crane calls. Unplanned downtime fell 42%, and maintenance costs dropped 27%, saving an estimated \$640K in the first year. Fleet availability climbed from 94.6% to 97.8%, strengthening the availability documentation the company files to justify contract pricing. Truck rolls on healthy machines fell by a third, returning technician hours to productive work. The prior year's 34-day outage, replayed against today's models, would have carried a three-week warning.

METRIC	BEFORE	AFTER
Failure warning lead	-	up to 3 weeks
Unplanned downtime	-	-42%
Maintenance costs	-	-27%
Fleet availability	94.6%	97.8%

The way the team worked changed as much as what the numbers said. The maintenance organization now plans in windows rather than reacting in emergencies, and the crane mobilization that once arrived as a cost shock is now a scheduled line item. The work-order history that made the models possible has become a self-reinforcing asset — every confirmed alert enriches the labeled data that sharpens the next generation of models. Deloitte's finding that mature predictive maintenance programs typically cut breakdowns by around 70% now reads internally as a trajectory rather than a benchmark: the fleet's alert-to-confirmation rate improved in every quarter of the first year. And the availability story told to regulators changed character, from documenting outages to evidencing a system that sees them coming.

08 What Your Business Should Do

Cascadia Ridge's playbook applies to any asset-intensive operator, and its first rule is sequencing. Instrument the assets that hurt most: rank assets by failure cost and downtime impact, then sensor the worst offenders first, because full-fleet coverage is a phase-three goal, not a sensible day-one requirement. Historize every work order now — start capturing clean, standardized work orders today so your models have training data six months from now, since failure models are only ever as good as their labeled history. Choose sensors against failure modes, working from OEM documentation rather than installing generic sensor kits blindly, because bearing wear, imbalance, overheating and oil degradation each demand different measurements. Then pair models with maintenance planners: wire alerts into your CMMS so forecasts become scheduled work orders with parts and crews assigned, because a prediction without a planning workflow is a novelty.

Operate the system for credibility. Set thresholds by severity so crews trust it — a few high-precision warnings beat constant noise, and every false alarm spends credibility you will need for the big catch. Compare suspects against their twins, using fleet-wide similarity views to separate true anomalies from units that simply run hot or sit in harsher terrain; contextual baselines cut false positives dramatically compared with fixed thresholds. Bundle alerts into planned windows, grouping nearby or related repairs to share crane and travel costs — the real economics of predictive maintenance live entirely in this bundling step. And report availability to the business, translating sensor alerts into availability, generation and dollar terms for executives, because programs funded on avoided-outage dollars survive budget cycles that dashboards full of model-accuracy metrics do not. The benchmarks justify the effort: McKinsey Global Institute estimates IoT-based predictive maintenance can cut downtime by up to 50% and maintenance costs by 10-40%, and US Department of Energy guidance reports

30-40% savings over reactive programs.

Before the first sensor kit ships, sit with the questions Cascadia Ridge's leadership had to answer honestly. What did your last three unplanned equipment failures cost — repairs, downtime and lost production combined? How much warning does your maintenance team get today before a critical asset fails? Which machines would you instrument first if you could only afford sensors on ten assets? Do you have clean failure history, or are work orders scattered across spreadsheets and memory? Would your planners trust a model's warning, and what evidence would it take to convince them? How often do technicians currently drive to a site and find nothing wrong? Take these to your next operations review with the action checklist in this practice book, and revisit them each quarter as your fleet ages and your alert library grows.

REAL-WORLD PROOF & BENCHMARKS

McKinsey Global Institute analysis estimates that IoT-based predictive maintenance can reduce equipment downtime by up to 50% and cut maintenance costs by 10-40%. - McKinsey Global Institute

Deloitte analysis reports that predictive maintenance programs typically reduce breakdowns by around 70%, increase productivity by 25% and lower overall maintenance costs by 25%. - Deloitte

US Department of Energy guidance on maintenance best practices reports that a functioning predictive maintenance program can yield savings of 30-40% over reactive approaches. - US DOE

THE PRACTICE PAGES

The 8-action checklist

Work through the list with your team. Tick each box as your organization completes the step, and fill in an owner and a target date so every action has a name and a deadline.

☐

1. Instrument the assets that hurt most

Rank assets by failure cost and downtime impact, then sensor the worst offenders first. Full-fleet coverage is a phase-three goal, not a sensible day-one starting requirement.

OWNER

TARGET DATE

STATUS

--	--	--

☐

2. Historize every work order now

Failure models are only ever as good as labeled history. Start capturing clean work orders today so your models have training data six months from now.

OWNER

TARGET DATE

STATUS

--	--	--

☐

3. Choose sensors against failure modes

Bearing wear, imbalance, overheating and oil degradation each demand different measurements. Work from OEM failure-mode documentation instead of installing generic sensor kits on every machine blindly.

OWNER

TARGET DATE

STATUS

--	--	--

☐

4. Pair models with maintenance planners

A prediction without a planning workflow is just a novelty. Wire alerts into your CMMS so forecasts become scheduled work orders with parts and crews assigned.

OWNER

TARGET DATE

STATUS

--	--	--

☐

5. Set thresholds by severity

Tune alerting so crews trust it: a few high-precision warnings beat constant noise. Every false alarm spends credibility you will need for the big catch.

OWNER

TARGET DATE

STATUS

--	--	--



6. Compare suspects against their twins

Fleet-wide similarity views separate true anomalies from units that simply run hot or sit in harsher terrain. Contextual baselines cut false positives dramatically compared with fixed thresholds.

OWNER

TARGET DATE

STATUS

--	--	--



7. Bundle alerts into planned windows

Group nearby or related repairs into scheduled visits to share crane and travel costs. The real economics of predictive maintenance live entirely in this bundling step.

OWNER

TARGET DATE

STATUS

--	--	--



8. Report availability to the business

Translate sensor alerts into availability, generation and dollar terms for executives. Programs funded on avoided-outage dollars survive budget cycles that dashboards full of model-accuracy metrics do not.

OWNER

TARGET DATE

STATUS

--	--	--

AFTER THE CHECKLIST: OUR TOP 3 PRIORITIES

From the eight actions above, write down the three you will start first - and the single number you expect each one to move. Three priorities with owners beat eight without any.

PRIORITY 1 NUMBER IT SHOULD MOVE	PRIORITY 2 NUMBER IT SHOULD MOVE	PRIORITY 3 NUMBER IT SHOULD MOVE
--	--	--

THE PRACTICE PAGES

Reflections for your team

Use these questions in your next planning session and type short answers into the boxes. The gaps you find are usually the fastest wins.

Q1. What did your last three unplanned equipment failures cost — repairs, downtime and lost production combined?

Q2. How much warning does your maintenance team get today before a critical asset fails?

Q3. Which machines would you instrument first if you could only afford sensors on ten assets?

Q4. Do you have clean failure history, or are work orders scattered across spreadsheets and memory?

Q5. Would your planners trust a model's warning, and what evidence would it take to convince them?

Q6. How often do technicians currently drive to a site and find nothing wrong?

THE PRACTICE PAGES

My 90-day plan

Finish the session by committing to a plan. Three phases, one goal each, one number that tells you it worked. Bring this page to your next leadership review - or to a free consultation with Emerson Technology, and we will pressure-test it with you.

Days 1-30 Focus: audit and alignment

PRIMARY GOAL**KEY ACTIONS****HOW WE MEASURE SUCCESS****OWNER**

Days 31-60 Focus: build and prove

PRIMARY GOAL**KEY ACTIONS****HOW WE MEASURE SUCCESS****OWNER**

Days 61-90 Focus: scale and measure

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Tip: save a copy per team so each group keeps its own answers. Revisit the reflections every quarter - the questions stay the same even as the numbers move.

BOOK THE REVIEW BEFORE YOU CLOSE THE BOOK

Plans without a review date drift. Before you close this practice book, agree on when you will come back to these pages, who will run that session, and when you will check the numbers again. If you want an outside pair of eyes, bring this page to a free consultation with Emerson Technology and we will pressure-test the plan with you.

REVIEW DATE	REVIEWED BY	NEXT CHECK-IN
------------------------------------	------------------------------------	--------------------------------------

Ready for a transformation like Cascadia Ridge Energy's?

Emerson Technology designs and builds the data, AI, cloud and automation systems behind results like these. On a free consultation call we will map this exact playbook to your business - your data, your constraints, your numbers.

Book a free consultation

support@emersontechnology.com

or start online: emersontechnology.com/contact