NB INDUSTRIES + SORBA.ai

SMART MACHINES. SIMPLE AI.

BEYOND MAINTENANCE: Using Predictive Data to Engineer Safer Asset Lifecycles

TABLE OF CONTENTS

EXECUTIVE SUMMARY

In modern mining operations, maintenance isn't just about uptime — it's about safety. This paper explores how predictive diagnostics, enabled by AI and real-time data, are redefining how mine sites extend equipment life, reduce unplanned failures, and engineer risk out of operations. Drawing on SORBA.ai's industrial applications and mining case insights, we show how a data-first approach enhances both operational efficiency and workforce wellbeing.

Extensive research by Ron Moore found a direct link between reactive maintenance and workplace injury rates. In a multi-year study of a large chemical facility, Moore reported a Pearson correlation coefficient of 0.83 between the number of corrective work orders and total injuries per year—highlighting that as reliability improves, safety outcomes also improve.

INTRODUCTION: THE SAFETY STAKES OF MAINTENANCE

In high-intensity mining environments, equipment failure is never just an operational inconvenience—it's a safety risk. The repercussions of a failed haul truck component, overheating engine, or blocked filter can cascade quickly, threatening not only production schedules but also the wellbeing of frontline workers.

Historically, maintenance strategies have been designed to keep machines running—not necessarily to keep people safe. Reactive repairs respond only after breakdowns occur. Scheduled servicing, while proactive in nature, often lacks precision—resulting in either premature interventions or missed warning signs.

But the ground is shifting.

With the rise of artificial intelligence, sensor networks, and real-time data streaming, mining operators now have the tools to move beyond traditional maintenance paradigms. Predictive diagnostics offer a smarter, more targeted approach—one that prioritises both asset performance and human safety by detecting early signs of failure long before they become incidents.

This white paper explores how predictive maintenance, powered by platforms like SORBA.ai, is helping mining operations extend the lifecycle of their critical equipment while actively reducing operational risk. We will examine real-world applications of condition monitoring, anomaly detection, and automated data cleansing—showing how they come together to form a safety-first, intelligence-driven maintenance strategy.

For forward-thinking mining leaders, the goal is no longer just uptime—it's safer uptime. And predictive data is the lever to make it happen.

THE HIDDEN RISK IN EVERYDAY FAILURES

Studies show a strong correlation between equipment reliability and workforce safety.

For example, data from a large chemical plant demonstrated a Pearson correlation coefficient (R) of 0.83 between the number of corrective/reactive work orders and total injuries per year — indicating that higher reliability directly contributes to lower injury rates.

This finding, reported by Ron Moore, reinforces the premise that a reliable plant is not just more efficient—it's inherently safer. (Source: https://www.lce.com/resources/a-reliable-plant-is-a-safe-plant-is-a-cost-effective-plant/)

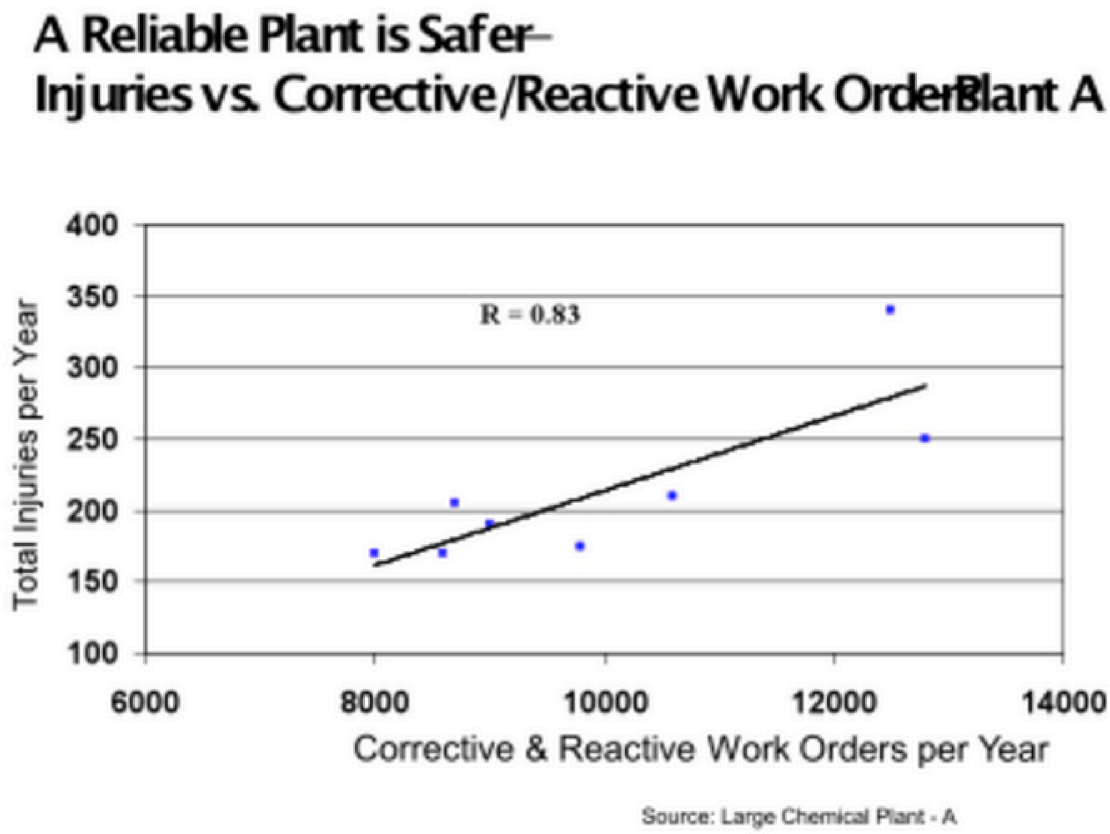

In mining, not all failures start with loud alarms or visible damage. Many of the most dangerous incidents begin quietly—through seemingly minor faults that go undetected during regular operations. When left unaddressed, these small anomalies accumulate stress, compromise structural integrity, and eventually lead to catastrophic equipment failures.

Take, for example, a partially obstructed engine air filter. At a glance, it may not trigger an immediate performance issue. But over time, that obstruction can lead to reduced airflow, inefficient combustion, and overheating—conditions that increase the risk of thermal fatigue and engine fire. Similarly, undetected air induction leaks allow dust particles to infiltrate engine systems, accelerating component wear and increasing the likelihood of ignition.

The reality is this: equipment doesn't just fail—it deteriorates. And deterioration leaves a trail of clues long before failure occurs.

These clues are detectable in the form of:

- Abnormal temperature rises
- Shifts in fluid quality or viscosity
- Recurring fault codes
- Changes in vibration signatures
- Declining system responsiveness

However, traditional maintenance systems—reliant on periodic inspections and reactive diagnostics—are ill-equipped to capture these early signals at scale. Technicians and supervisors are often left responding to symptoms rather than root causes, increasing the frequency of breakdowns and the risk of workplace injuries.

Moreover, unplanned maintenance can introduce further risk. Emergency repairs are often conducted under pressure, in harsh environments, and sometimes in unsafe conditions. The likelihood of human error increases, and the stress placed on frontline personnel is amplified.

By contrast, predictive diagnostics enable maintenance teams to catch problems upstream—before they escalate into safety events. This shift transforms maintenance from a necessary operational cost into a strategic safety control.

In mining, where the margin for error is narrow and the consequences severe, identifying the hidden risks behind everyday mechanical behaviours is no longer optional—it's essential.

PREDICTIVE MAINTENANCE TECHNOLOGIES: WHAT'S NOW POSSIBLE

The evolution of predictive maintenance is not a story of theory—it's a story of technology catching up with the needs of frontline operations.

Until recently, maintenance in mining was constrained by two extremes: reactive fixes that came too late, and time-based schedules that came too often. Both approaches are inefficient, and neither is aligned with actual asset condition. What's changed is the ability to continuously monitor, interpret, and act on equipment data in real time—at scale.

The Building Blocks of Predictive Maintenance

Mining leaders must rethink how they introduce and champion AI systems. Building trust means:

Sensor Networks and IoT Integration

Modern mobile and fixed mining assets are equipped with a growing range of onboard sensors. These collect vital condition data such as:

- Temperature
- Pressure
- Vibration
- Oil quality
- Load balance
- RPM and duty cycles

With Industrial IoT (IIoT) systems in place, this data is streamed from field assets to a central repository or edge gateway, creating a live digital heartbeat of every machine.

Machine Learning and Asset Modelling

This is where raw data becomes insight.

Machine learning algorithms are trained on historical and real-time asset behaviour to understand what "normal" looks like. Once a baseline is established, the system can detect anomalies—those subtle deviations that precede mechanical failure. Over time, the models learn, refine, and adapt to each asset's unique operating context.

Key outputs include:

- Early fault detection
- Remaining useful life estimation
- Failure probability scoring
- Maintenance prioritisation

Real-Time Dashboards and Alerts

Information is only useful when it's actionable. Predictive platforms now provide dashboards that visualise asset health status, trending anomalies, and recommended interventions. Custom alerts—delivered via email, SCADA, or even direct to handheld devices—allow maintenance planners to act decisively before issues become incidents.

Closed-Loop Feedback into Operations

Advanced predictive systems can go a step further by feeding insights directly into operational logic. For example, if vibration exceeds threshold on a haul truck's chassis, a signal can be sent to adjust speed or route in real time—protecting the asset while maintaining production continuity.

This kind of closed-loop optimisation is crucial in complex, autonomous, or remote mining environments, where fast decision-making can prevent system-wide disruption.

From Potential to Practice

These technologies are no longer confined to pilot projects or academic papers. They're being deployed in working mine sites today—often without the need to overhaul existing infrastructure.

By aligning machine learning with condition monitoring, mining operators now have the power to identify risks days or weeks in advance, avoid catastrophic failures, and shift from firefighting mode into intelligent, proactive planning.
The result? Safer equipment, less downtime, and a maintenance culture that leads—not lags—operational strategy.

WHAT SORBA.AI OFFERS MINING OPERATIONS

Implementing predictive maintenance in mining isn't just about collecting data—it's about making that data meaningful, actionable, and scalable. That's where SORBA.ai stands out.

Built specifically for industrial environments, SORBA.ai is an AI-driven platform designed to help mining operations monitor, model, and manage equipment health in real time. It bridges the gap between raw sensor data and informed maintenance decisions—without requiring a major overhaul of existing infrastructure.

A. Seamless Integration with Existing Systems

Mining operations don't have time to rip and replace. SORBA.ai integrates directly with:

- SCADA platforms like Ignition
- Industrial communication layers like Kepware
- Data historians and PLC-based systems

This allows SORBA to ingest asset data natively—without custom code, without rewriting PLC logic, and without interfering with production.

B. Real-Time Monitoring at the Edge

SORBA supports edge computing, enabling real-time processing and decision-making directly at the source of data—whether that's a haul truck, processing plant, or pumping station. This reduces latency, improves fault response time, and ensures reliability in remote or low-connectivity environments.

C. Predictive Model Training for Every Asset

Every piece of mining equipment behaves differently depending on duty cycle, environment, and age. SORBA trains bespoke machine learning models for each asset using:

- Live streaming data
- Cleaned historical datasets
- Feedback from anomalies and outcomes

This results in accurate fault detection tailored to actual site conditions—not generic OEM guidelines.

D. Automated Data Cleansing

Mining data is messy. From sensor drift to communication dropouts, poor-quality data can sabotage even the best algorithms.

SORBA's built-in automated data cleansing engine filters out:

- Spikes and dropouts
- Duplicate readings
- Null or frozen values

This ensures models are trained on clean, consistent data—improving accuracy and reducing false positives.

E. Anomaly Detection and Risk Prioritisation

Once trained, SORBA continuously scans incoming data for deviations from normal behaviour. These anomalies are:

- Scored by risk level
- Categorised by fault type (e.g. vibration, temperature, pressure)

- Linked to probable root causes
- Routed to relevant personnel for review

This gives maintenance teams a clear, prioritised task list—cutting through the noise and focusing effort where it matters most.

F. Closed-Loop Feedback and Control

In advanced deployments, SORBA can issue commands directly to control systems or operator dashboards. For example:

- Trigger derating procedures on overheating equipment
- Recommend route adjustments for fatigue-prone haul trucks
- Adjust servicing schedules based on usage and wear

This closed-loop optimisation turns predictive insight into operational action—reducing downtime and protecting equipment in real time.

G. Built for Harsh, Heavy Industries

SORBA.ai was built with the realities of mining in mind:

- Works in noisy, high-vibration, dusty environments
- Handles large-scale deployments across fleets and plants
- Designed for minimal maintenance and operator involvement

In short: it's a platform made for mining—not retrofitted from office-based analytics.

SAFER LIFECYCLES THROUGH INTELLIGENT SCHEDULING

In mining, maintenance isn't just about fixing what breaks—it's about preventing failure before it becomes a threat to people and productivity. Predictive maintenance enabled by platforms like SORBA.ai doesn't just detect risks earlier—it changes how maintenance is scheduled, how resources are allocated, and how safe a worksite can be.

A. From Reactive Repairs to Risk-Based Planning

Traditional maintenance is often either:

- Reactive: responding after something breaks, or
- Time-based: servicing equipment on a fixed calendar, regardless of condition.

Both approaches can create hidden risks. Reactive maintenance typically occurs under pressure, increasing the likelihood of rushed decisions and unsafe working conditions. Time-based servicing can lead to unnecessary intervention—or worse, missed deterioration in-between cycles.

Predictive maintenance offers a third, smarter path: condition-based scheduling.

With real-time monitoring and AI-driven insights, maintenance teams can now:

- Intervene before a part fails—not after
- Extend the service interval of healthy components
- Prioritise interventions based on actual risk, not just time elapsed

Moore's analysis of industrial operations also shows that organisations with fewer reactive work orders see significantly lower injury rates — evidence that intelligent scheduling reduces not only mechanical failures but human harm.
(Source: https://www.lce.com/resources/a-reliable-plant-is-a-safe-plant-is-a-cost-effective-plant/)

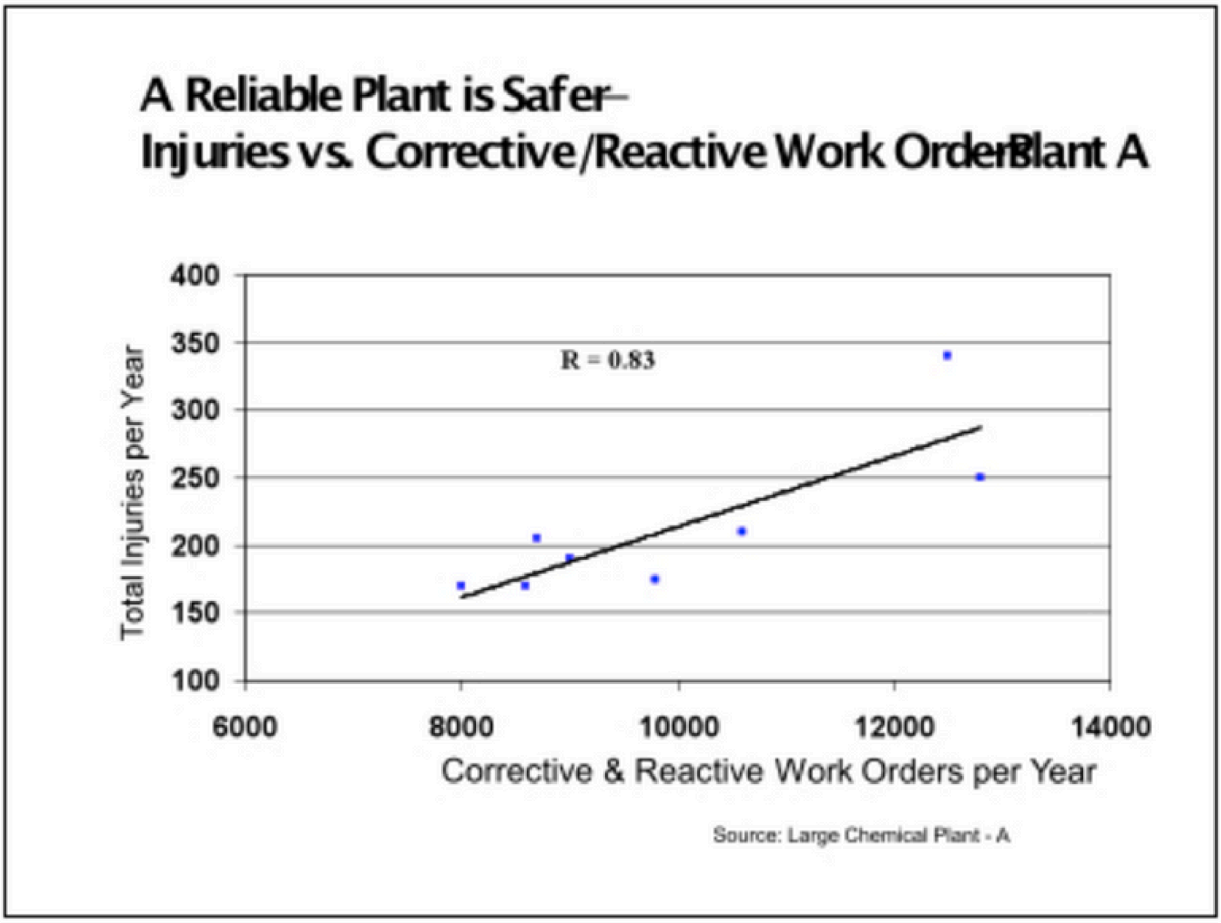

B. Avoiding Dangerous Scenarios Before They Unfold

Early detection of problems allows mines to prevent the most dangerous failure scenarios.

For example:

Condition	Predictive Detection	Prevented Outcome
High particulate levels in engine bay	Airflow and pressure anomalies	Fire risk due to dust ignition
Hydraulic system instability	Pressure and flow monitoring	Sudden line failure or fluid leaks
Chassis fatigue on haul truck	Suspension trend analysis	Frame cracking or catastrophic collapse

These aren't just equipment problems—they're safety events waiting to happen.

C. Building Operator Confidence and Reducing Stress

An often overlooked benefit of predictive maintenance is the psychological safety it provides.

When operators know their equipment is being continuously monitored and maintained based on real conditions—not guesswork—they gain confidence in the machinery they rely on daily. It reduces second-guessing, builds trust in alerts, and fosters a more proactive mindset across the site.

This is especially important in autonomous or semi-autonomous operations, where machine trustworthiness is critical.

D. Improving Maintenance Team Morale and Workflow

Predictive maintenance helps maintenance teams shift from “firefighting mode” to “planning mode.” Instead of scrambling to respond to breakdowns, they can:

- Organise workdays around prioritised alerts
- Allocate resources more effectively
- Reduce the number of emergency callouts and out-of-hours servicing

This leads to improved morale, clearer communication, and safer execution of tasks.

E. A Step Toward Autonomous and Resilient Mining

As mining becomes more digitised and autonomous, condition-based maintenance becomes the foundation for safe operations at scale. AI doesn't replace maintenance teams—it equips them with tools to stay ahead of failure, eliminate guesswork, and maintain the integrity of increasingly complex systems.

SORBA.ai doesn't just support this evolution—it accelerates it.

CASE EXAMPLE: SORBA.AI IN ACTION

Case Study: Early Fault Detection in Caterpillar 797 Mining Truck Using SORBA.ai

Background:

The Caterpillar 797 is a heavy-duty mining truck equipped with various sensors to monitor engine performance and other critical operational parameters. In the mining industry, even minor issues can lead to substantial operational losses. Traditionally, maintenance teams rely on preset thresholds to trigger alarms, often resulting in reactive rather than proactive maintenance.

Scenario:

A mining company operating a Caterpillar 797 truck faced recurring exhaust gas temperature (EGT) alarms, leading to multiple shutdowns and disrupted operations. To implement a more proactive approach, the company sought a solution to detect issues before they caused significant disruptions.

Challenges:

The primary challenge was the limited ability of conventional monitoring systems to provide early warnings. The existing system only triggered alarms when parameters exceeded specific limits, often leaving insufficient time for preemptive action, which resulted in unplanned downtime and increased operational costs.

SORBA.ai's Approach:

To address these challenges, the company deployed SORBA.ai's predictive maintenance platform, leveraging its AI-driven Sorbots for an offline analysis of historical data from the truck's sensors.

Figure 1: Data received for CAT Engine Tags - 6 hours before the alarm with SORBA.ai

Implementation Steps:

Data Ingestion and Analysis: *SORBA.ai was fed with the truck's sensor data, including all engine tags. The AI model established a baseline of normal operational behaviour by analysing this historical data.*

Anomaly Detection: *The "Engine Supervisor" Sorbot detected a deviation from normal behaviour patterns six hours before the truck's onboard monitoring system triggered an EGT alarm. This early detection enabled the maintenance team to investigate the issue before it escalated into a more severe problem.*

Detailed Diagnostics: *A second Sorbot, known as a "Lackey," was configured to focus specifically on identifying intake leaks. This Sorbot successfully pinpointed an anomaly caused by a loose hose clamp in the intake system, providing a 24-hour advance warning before the truck's system alarmed.*

Tag Ranking and Precision: *SORBA.ai also utilised a feature known as "tag ranking" to prioritise data streams that exhibited the most significant deviations, directing the maintenance team to the intake manifold where the loose clamp was suspected.*

Results:

By employing SORBA.ai, the mining company achieved significant improvements in fault detection and maintenance efficiency. The early warnings provided by the Sorbots allowed the team to take preventive measures, reducing unplanned downtime and optimising resource allocation.

BEYOND COST: BUILDING THE SAFETY AND ROI CASE

When predictive maintenance is framed solely as a cost-saving initiative, its true value is often underestimated. Yes—fewer breakdowns, fewer parts, and fewer emergency repairs absolutely improve the bottom line. But in mining, the more transformative return comes from what predictive maintenance prevents: safety incidents, operational chaos, and damage to workforce morale.

A. The ROI of Safer Operations

In high-risk industrial environments, safety isn't a "soft" metric —it's a critical financial variable. Incidents caused by equipment failure can result in:

- Worker injury or fatality
- Lost-time incidents (LTIs)
- Regulatory penalties and shutdowns
- Damage to reputation and social license to operate

By catching faults before they escalate, predictive maintenance becomes a risk control mechanism—lowering incident frequency, severity, and associated costs.

B. Direct Financial Benefits

Mining operations using platforms like SORBA.ai typically see tangible cost reductions in:

- Unplanned downtime: Avoiding production halts from unexpected equipment failure
- Labour efficiency: Reducing the need for reactive/emergency maintenance crews
- Component life extension: Replacing parts based on condition, not calendar
- Optimised inventory: Less stockpiling of “just in case” parts
- Fewer out-of-hours repairs: Reduced fatigue-related errors and overtime costs

Combined, these outcomes result in higher asset availability, longer equipment life, and improved budget predictability.

C. Intangible—but Strategic—Gains

While hard-dollar savings are important, many of the most valuable outcomes are intangible:

Outcome	Strategic Value
Increased operator trust in equipment	Fewer human overrides, smoother production
Fewer emergency callouts	Improved wellbeing and retention of skilled staff
Better planning data	Smarter CAPEX decisions and resource allocation
ESG alignment	Demonstrates responsible, forward-thinking operations

Predictive maintenance contributes to broader Environmental, Social and Governance goals, particularly under safety, resource efficiency, and workforce health pillars. For companies reporting under sustainability frameworks, this data-driven approach is easy to justify and communicate to stakeholders.

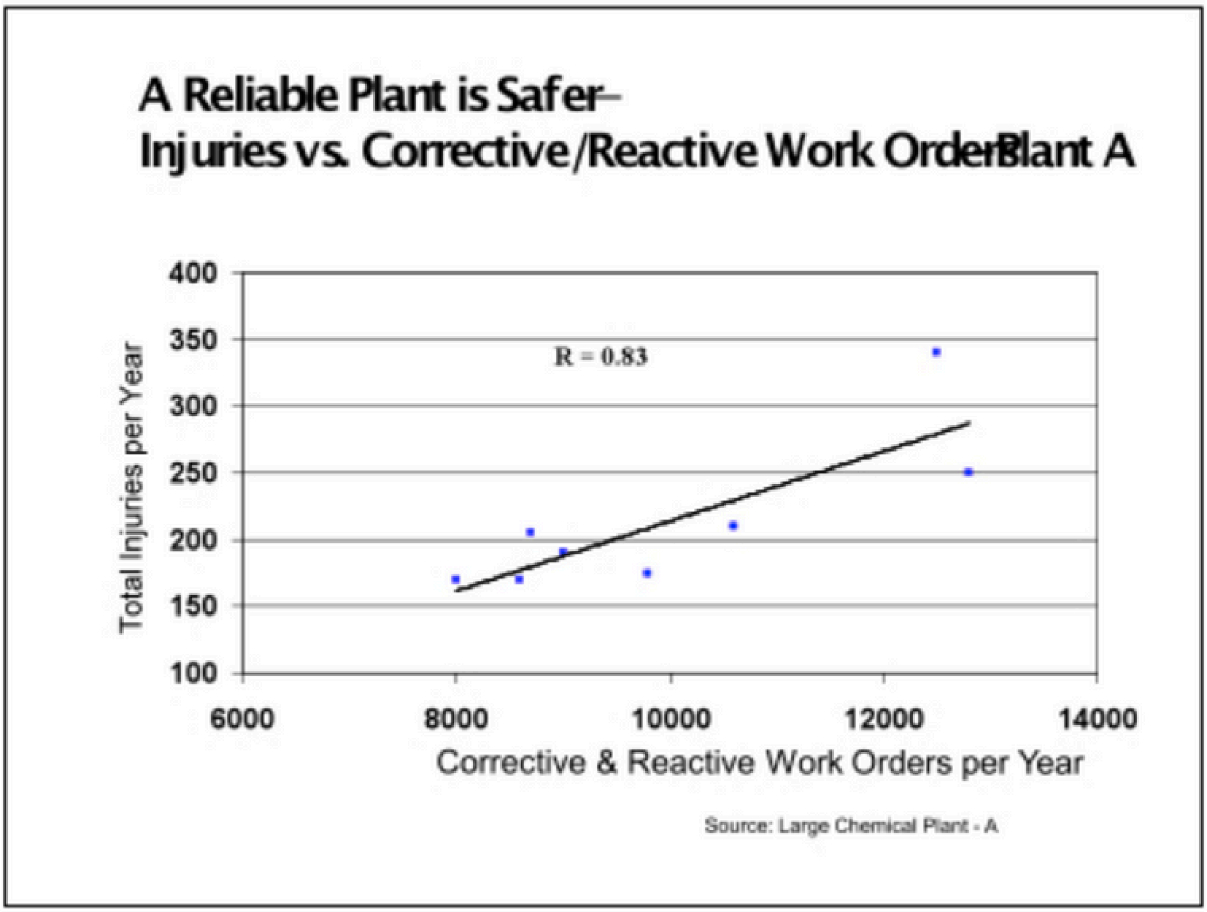

This is supported by findings from Ron Moore, who demonstrated a near-linear correlation (R=0.83) between reactive maintenance workload and injury rates. In practice, this means every step toward predictive reliability also reduces risk exposure and incident frequency.
(Source: https://www.lce.com/resources/a-reliable-plant-is-a-safe-plant-is-a-cost-effective-plant/)

D. Unlocking the Compounding Advantage

The earlier predictive systems are deployed, the greater the compound return. Every hour of clean data and every avoided fault improves the model's precision, expanding insight across:

- Asset classes
- Site conditions
- Operational strategies

The insights don't just improve individual machines—they enable smarter system-wide decision-making.

E. Making the Business Case Stick

To secure executive buy-in for predictive maintenance, industrial leaders should focus their case on:

- **Cost of inaction:** Quantify historical failures and near-misses that predictive maintenance could have prevented

- **Fast wins:** Identify one high-risk, high-cost asset to pilot—like haul trucks or crushers

- **Safety-first framing:** Align predictive maintenance with your site's safety KPIs and incident reduction targets

- **Scalable roadmap:** Present a rollout plan that builds credibility through measurable stages

GETTING STARTED WITH SORBA.AI

1. **Choose a High-Impact Asset:** Start with an asset or system that's safety-critical and data-rich—such as haul trucks, crushers, or conveyors.
2. **Conduct a Data Audit:** Assess what data is already available from your SCADA, PLC, or sensor network. Identify connectivity points through platforms like Ignition or Kepware.
3. **Clean and Prepare the Data:** SORBA.ai automatically cleans the data by removing spikes, gaps, and noise—ensuring models are built on accurate, high-quality inputs.
4. **Train Predictive Models:** Machine learning models are trained on your equipment's unique behaviour to establish baseline performance and detect early signs of failure.
5. **Configure Alerts and Dashboards:** Customise real-time alerts and visual dashboards to suit your site's maintenance workflows. Alerts are prioritised by risk level and routed to the right team.
6. **Act and Optimise with Feedback:** Maintenance teams respond to alerts and provide feedback on their accuracy. This input is used to refine model performance and reduce false positives.
7. **Scale Across Operations:** Once proven, replicate the setup across additional assets, fleets, or sites. The system is modular and easily scalable without disrupting existing infrastructure.
8. **Ongoing Support and Integration:** SORBA.ai provides training, model reviews, and integration with your maintenance systems to ensure continuous improvement and long-term success.

CONCLUSION

In today's mining environment, where safety is paramount and downtime is costly, relying solely on reactive or time-based maintenance is no longer good enough. The future of safe, efficient operations lies in predictive diagnostics—powered by clean data, advanced algorithms, and systems that learn with every hour of machine operation.

This white paper has shown how AI-driven platforms like SORBA.ai are helping mining leaders move beyond traditional maintenance. By identifying subtle signs of failure early, streamlining intervention, and supporting intelligent scheduling, predictive maintenance is shifting maintenance from a cost centre to a strategic safety control.

It's not just about longer equipment life—it's about fewer injuries, smarter decision-making, and building a culture where safety and performance go hand in hand.

For organisations aiming to lead in both production and protection, the path is clear: **Use data not just to maintain, but to prevent. Not just to optimise, but to protect.**

Predictive maintenance is no longer aspirational—it's operational. And with the right tools, the journey from reactive to proactive can begin today.

APPENDIX

A. Key Terms and Concepts

Predictive Maintenance (PdM): Maintenance strategy that uses real-time data and AI models to forecast equipment failure before it occurs.

Condition Monitoring: The continuous or periodic tracking of key operational parameters (e.g., vibration, temperature, pressure) to assess asset health.

Anomaly Detection: The process of identifying data patterns or equipment behaviours that deviate from the norm, signalling potential failure.

Closed-Loop Feedback: A system where AI outputs (e.g., alerts or control signals) are sent directly back to operations or control logic for automated action.

Edge Computing: Processing data close to its source (e.g., on a truck or at a remote substation) to enable fast response and reduce bandwidth requirements.

Data Cleansing: The removal of errors, duplicates, noise, or irrelevant data points from a dataset to improve AI model accuracy.

Fault Code Frequency: The rate at which diagnostic trouble codes are triggered, used as an indicator of system instability or degradation.

Thermal Stress: Physical strain on components caused by rapid or extreme temperature fluctuations, often leading to cracks or material fatigue.

B. Common Safety-Critical Faults in Mining Equipment

Fault Type	Potential Risk
Air Induction Leaks	Dust ingress → accelerated wear → fire hazard
Partially Obstructed Filters	Combustion instability → overheating → structural fatigue
Hydraulic Overpressure	Seal rupture → fluid leaks → burn or slip hazard
Excess Vibration	Component loosening → catastrophic failure
Oil Degradation	Loss of lubrication → friction wear → breakdown or fire

C. Additional Learning Resources

SORBA.ai University (White Papers, Webinars)
https://university.sorba.ai/

SORBA.ai Website (Case Studies)
https://www.sorba.ai/

Article from A Reliable Plant is a Safe Plant is a Cost-Effective Plant by By Ron Moore
https://www.lce.com/resources/a-reliable-plant-is-a-safe-plant-is-a-cost-effective-plant/

ABOUT THE AUTHOR

Peter Horsburgh is a renowned expert in Reliability Engineering with over two decades of experience in manufacturing. He is the author of the popular book '5 Habits of an Extraordinary Reliability Engineer' and has worked with leading global manufacturers to integrate AI and machine learning into quality control, driving improvements in predictive maintenance and process optimization.

Peter leads efforts to solve manufacturing challenges through advanced technologies. A sought-after speaker and educator, his work inspires professionals to adopt innovative strategies for enhancing reliability and operational excellence.

Peter Horsburgh

SORBA.ai APAC Managing Director
and Renowned Author

BOOK A CALL WITH ME!

+61 427 151660

peter@sorba.ai

www.sorba.ai