Zero Downtime:
From Closed Loop to Competitive Advantage

Part II

Sponsored by: AssetWatch

Modern reliability programs are increasingly built around a closed-loop model: **Detect > Diagnose > Action > Improve.** This model involves implementing a predictive maintenance (PdM) strategy through continuous monitoring, diagnostics, and structured interventions to identify developing faults early enough to prevent failures. The cone crusher example in Part 1, which avoided an estimated $240,000 in downtime, illustrated how early detection can positively affect financial and operational outcomes. This potential is significant, as PdM strategies have been shown to **reduce unplanned downtime by 30% to 50% in many industrial operations.**[1]

This leads to an important question. **Why do some plants implement a closed-loop model and plateau while others turn it into a measurable competitive advantage?**

The difference is not the technology itself. It's how the organization develops around it.

Many plants deploy condition monitoring solutions and successfully prevent failures. However, relatively few translate those successes into sustained performance improvements such as stable throughput, better capital utilization, and reduced operational risk. In many cases, reliability initiatives stop at system deployment rather than becoming integrated into decision-making and planning.

This paper examines what separates organizations that implement reliability tools from those that manage reliability at scale. The difference lies in building reliability as an organizational capability rather than treating it solely as a maintenance initiative.

Asset Performance Maturity Model

Reliability maturity does not develop from a single initiative. It emerges when multiple organizational capabilities improve together. The **Asset Performance Maturity Model** provides a framework for understanding how those capabilities build on one another and where organizations tend to stall. Instead of focusing only on technology or tools, the stack looks at the broader systems that determine whether reliability efforts translate into sustained performance.

The model is built around four organizational layers that together determine how effectively reliability is executed and scaled. Each layer supports the one above it, and when one layer is weak, it limits the impact of the others.

The table below identifies the core question each layer answers, what it controls inside the organization, and what typically happens when that layer is either weak or mature.

Layer	Core Question	What it Controls	If Weak	If Mature
Technical Capability	Are we detecting and diagnosing risk early?	Visibility and fault lead time	Late detection, reactive maintenance, hidden risk	Predictive control, quantified asset risk, extended P-F lead time
Process Discipline	Do insights consistently convert into planned action?	Execution stability and response quality	Alert fatigue, inconsistent follow-through, emergency work	Engineered intervention, defined SLAs, optimal intervention timing
Financial Integration	Is reliability translated into financial outcomes?	Investment influence and capital allocation	Anecdotal ROI, maintenance seen as cost center	Finance-validated savings, capital deferral, working capital optimization
Cultural Ownership	Is reliability shared across the enterprise?	Sustainability, scale, and governance	Initiative fatigue, siloed accountability	Enterprise risk alignment, shared KPIs, executive-level visibility

This structure reinforces an important principle.

Technical capability alone does not create competitive advantage.

↓

Strong diagnostics without disciplined execution still produce variability.

↓

Disciplined execution without financial validation limits influence.

↓

Financial validation without cultural ownership fails to sustain momentum.

When all four layers mature together, reliability evolves from a maintenance function into an enterprise risk management capability.

Technical Capability: From Monitoring to Risk Intelligence

Part I established the foundation of detection and diagnostics where continuous monitoring and structured analysis reduce blind spots to extend the P-F window. That foundation is essential, but maturity requires moving beyond visibility into risk intelligence.

The shift begins with prioritization. Organizations must evaluate whether:

- Insights are ranked by true level of severity rather than alarm volume
- Asset criticality is factored into decision making
- Leading indicators are replacing lagging metrics such as downtime reports
- Teams are tying fault detection, resolution, and precision maintenance to production impact

Without this evolution, teams measure activity. With it, they actively manage operational uncertainty.

In advanced environments, the focus shifts from simply detecting problems to understanding how early risks are identified and how effectively they are managed. This requires metrics that measure risk reduction and fault visibility rather than just system coverage.

Key indicators typically include:

- **A live Asset Risk Index,** which ranks assets based on the likelihood of failure.

This allows teams to prioritize intervention based on business impact rather than alarm volume.

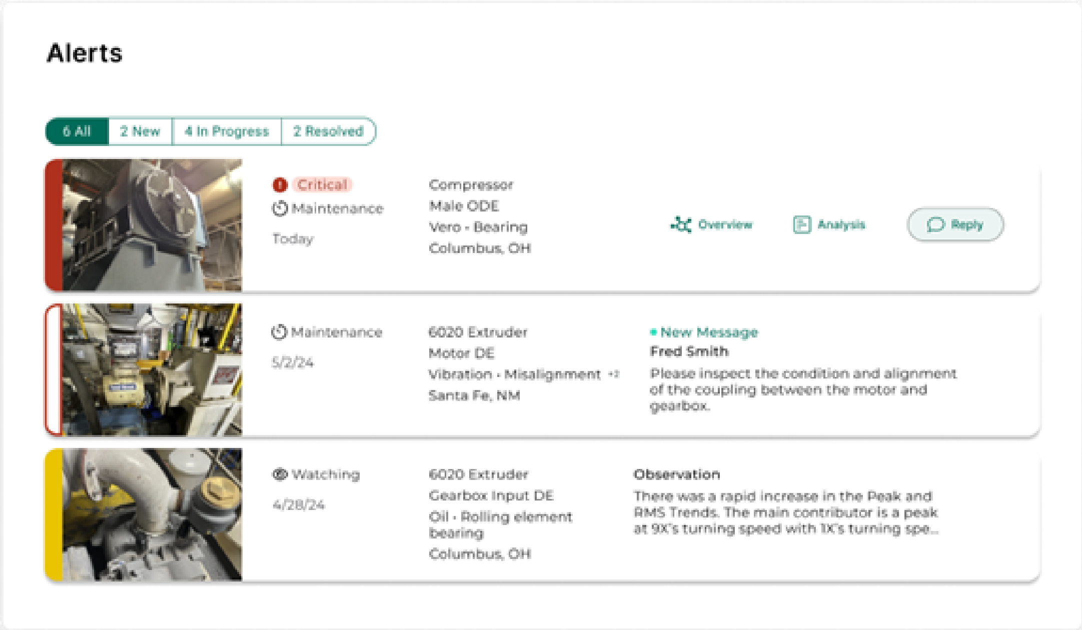

- **Mean Time to Comment (MTTC)** measures the time from alert creation to the first meaningful human acknowledgment. This serves as a proxy for product adoption and engagement, indicating whether teams are actively monitoring, trusting, and interacting with the system in a timely manner.
- **Mean Time to Resolve (MTTR)** measures the time from alert creation to confirmed resolution or closure. This reflects operational effectiveness by showing how efficiently teams move from detection to action and ultimately eliminate the issue.
- **Percentage of failure modes detected within the P-F window,** which measures how often faults are identified early enough for planned intervention rather than reactive repair. This metric is critical because many failures are still addressed reactively. **Studies show that about 40% of emergency maintenance calls occur after hours and 25% occur on weekends,** highlighting how often issues escalate outside of planned maintenance windows when faults are not detected early enough.[2]
- **Early fault detection lead time trending,** which tracks how much advance warning is gained before failures or severe degradation occur.

Together, these indicators shift the focus from coverage to control. Instead of asking how many assets are monitored, organizations should evaluate how much operational risk has been reduced, how early degradation is being detected, and whether the system is consistently increasing the time available to intervene before failure occurs.

Process Discipline: From Alerts to Engineered Outcomes

A closed-loop system may exist. The real question is whether it is institutionalized.

Detection without disciplined execution creates noise. Alerts are generated, work orders are opened, but responses vary by shift, supervisor, or workload. Maturity at this layer is defined by operational rigor.

Organizations strengthen process discipline by implementing:

- Standard intervention playbooks for common fault types
- Cross-functional reliability reviews
- Defined timelines to react based on criticality
- Escalation logic tied directly to quantified risk

These structures ensure that similar faults trigger consistent responses. They reduce variability and remove ambiguity from decision making.

In advanced environments, execution quality is measured, not assumed. Generating alerts is only part of the equation. What matters is how consistently those alerts lead to timely and effective intervention. Organizations that reach higher maturity levels track specific indicators that show whether predictive insights are translating into reliable operational outcomes.

Key indicators include:

- **A defined Alert-to-Intervention Service Level Agreement,** which establishes how quickly different types of faults must be evaluated and addressed based on severity and risk. This prevents alerts from sitting unresolved while work is scheduled reactively.
- **Percentage of predictive alerts executed within the optimal intervention window,** which measures how often maintenance actions occur early enough to avoid production disruption while still maximizing component life.
- **Reduction in repeat defect rates,** which indicates whether corrective actions are addressing root causes rather than repeatedly repairing the same issues.

- **Monthly trending of planned versus unplanned work,** which reflects whether predictive insights are increasing the share of maintenance that can be scheduled and controlled.

For example, two plants may detect the same fault at the same time. In one facility, the alert is acknowledged but intervention is delayed due to competing priorities, and the issue eventually disrupts production. In the other, the fault falls into a predefined response category and is scheduled within an established intervention window. The repair occurs during planned downtime, and production impact is minimal.

Modern metrics are critical to evaluating the success of your program.

Read about those metrics here.

The difference is not the alert itself. It is the discipline behind the response. When process discipline is strong, the closed loop produces predictable outcomes rather than inconsistent reactions.

This consistency extends beyond maintenance performance and begins to influence broader business outcomes. Reliable execution improves customer satisfaction by reducing unexpected disruptions, supports on-time-in-full (OTIF) delivery performance, and gives sales teams greater confidence to commit to production schedules without the risk of unplanned downtime.

Financial Integration: From Avoided Downtime to EBITDA Impact

This is the most significant leap in maturity.

In Part I, we highlighted a $240,000 avoided failure in the cone crusher case, demonstrating the value of predictive reliability. However, a single win does not create structural advantage. Maturity requires translating reliability into measurable financial impact.

Key questions at this layer include:

- Is reliability reported in operational terms or financial terms?
- Does Finance validate avoided cost?
- Is reliability influencing capital planning?

In advanced environments, financial integration looks like:

- **Verified Avoided Downtime that is finance-validated.** When predictive insights prevent a failure or major disruption, the avoided production loss, repair costs, and operational impacts are formally documented and reviewed with finance. This validation ensures savings are credible and consistently measured rather than estimated by maintenance alone. It also captures broader financial effects such as investments in workforce upskilling, adjustments to insurance premiums tied to risk reduction, and improvements in spare inventory levels and working capital efficiency.
- **Tracking improvements in working capital and spare inventory.** As reliability improves, organizations often reduce the need for excess spare parts, emergency purchases, and safety stock. **In fact, predictive maintenance programs can reduce spare inventory requirements by 15% to 20%,** freeing up capital that would otherwise remain tied up in parts inventory while still maintaining confidence in asset availability.[2]
- **Production stability and revenue protection.** Instead of relying on traditional equipment effectiveness metrics, mature programs connect reliability improvements directly to stable production and protected revenue. Early fault detection and timely intervention reduce unplanned disruptions, stabilize throughput, and ensure committed production targets are met.

- **Enterprise-wide trending of asset risk exposure.** Rather than focusing only on past failures, organizations track how overall risk across the asset base is changing over time. This allows leadership to see whether reliability initiatives are actually reducing the likelihood and potential impact of future disruptions.

For example, instead of simply reporting that a failure was prevented, an optimal organization quantifies the validated cost avoided, the throughput protected, and the capital deferred. Finance recognizes the impact, and leadership incorporates reliability metrics into forward planning.

ROI on PdM implementations averages 10x within 2 years.[3]

At this stage, reliability moves beyond preventing failures. It becomes a financial control mechanism that protects margin, reduces variability, strengthens long-term competitiveness, and can even start transitioning reliability and maintenance teams to cost saving centers.

Cultural Ownership: From Maintenance Tool to Enterprise Strategy

Part I emphasized the partnership between human expertise and AI. Cultural ownership takes reliability a step further, elevating it from a departmental initiative to an enterprise strategy. Even the best technical systems and financial validation cannot sustain long-term impact without it.

At this level, reliability becomes part of how the organization thinks and operates. Operations must trust the data, leadership should review asset risk regularly, and technicians need structured root-cause feedback to continuously improve processes. Reliability also begins to influence broader strategic planning, including quarterly performance reviews.

The goal is to move towards an organization where cultural ownership is formalized. An enterprise reliability council may guide shared KPI ownership across Operations, Maintenance, and Finance. Zero-unplanned-failure targets are set for critical assets, and reliability metrics are embedded into performance evaluations.

The Five Stages of Reliability Maturity

Maturity progression typically follows a series of stages, each representing a higher level of visibility, discipline, and organizational alignment. The framework below outlines five stages of reliability maturity, illustrating how organizations move from reactive maintenance toward enterprise-level reliability that supports operational stability and long-term performance.

Stage 1

In Stage 1, failures define priorities and downtime drives urgency, leaving little visibility into fault lead time or opportunities for early intervention.

Dimension	Description
Leadership Sees	Downtime reports, overtime spend, and budget pressure. Reliability is viewed primarily as a maintenance expense.
Technicians Experience	Firefighting, overtime, emergency work, and constant schedule disruption.
Metrics That Dominate	Unplanned downtime, maintenance spend, Mean Time Between Failure (MTBF).
What Unlocks Stage 2	Deployment of monitoring systems and basic fault detection capability to create early visibility.

Across the UK, US, and Germany, new data shows 46% of manufacturers reported between six and ten downtime incidents weekly, while for 15% the figure is between 11 and 20.

Nearly half (45%) said outages last up to 12 hours, with 17% reporting that incidents had stretched to 72 hours.[4]

Stage 2

Sensors are deployed and alerts increase, but diagnostics remain inconsistent, and prioritization is still developing, leaving ROI largely anecdotal rather than validated.

Dimension	Description
Leadership Sees	Increased data visibility and activity, but unclear or inconsistent business impact.
Technicians Experience	Frustration due to increased alerts and work orders, in addition to firefighting.
Metrics That Dominate	Alert counts, Mean Time to Comment (MTTC), overtime spend.
What Unlocks Stage 3	Standardized intervention processes and documented, repeatable financial wins tied to predictive action.

Stage 3

Alerts consistently convert into planned work. Early financial wins are documented, and reliability reviews begin to formalize cross-functional alignment.

Dimension	Description
Leadership Sees	Fewer emergencies, improved schedule stability, more predictable execution, and improved metrics.
Technicians Experience	Clearer prioritization, structured planning, and a shift from reactive to planned interventions.
Metrics That Dominate	Mean Time to Resolve (MTTR), reduced unplanned downtime, reduced line stoppages.
What Unlocks Stage 4	Finance validation of savings and implementation of risk-weighted asset prioritization.

Organizations adopting predictive maintenance often reduce maintenance costs by 10–40%.[3]

Stage 4
Finance validates savings. Asset prioritization becomes risk-weighted, and reliability begins influencing budget planning and capital allocation decisions.

Dimension	Description
Leadership Sees	Quantified savings, improved capital discipline, and increased executive visibility into asset risk.
Technicians Experience	Clear alignment between technical work and measurable business impact.
Metrics That Dominate	Production hours protected (per month/quarter).
What Unlocks Stage 5	Formal enterprise governance and zero unplanned failure targets for critical asset classes.

Stage 5
Reduced downtime is achieved for critical asset classes. Reliability drives throughput stability, and predictive metrics influence capacity planning and long-term strategy.

Dimension	Description
Leadership Sees	Margin protection, throughput confidence, reduced operational volatility, and reliability embedded in strategic planning.
Technicians Experience	Fewer emergencies, higher precision work, and strong ownership of asset health.
Metrics That Dominate	Risk reduction trends, throughput stability, and enterprise avoided downtime.
What Sustains This Stage	Continuous system learning, cross-site benchmarking, and deeply embedded cultural ownership.

These five stages illustrate that reliability maturity develops through progression rather than a single implementation. Many organizations deploy monitoring technologies but remain in the early stages because processes, financial alignment, and organizational ownership have not matured alongside the technology. As each stage builds on the one before it, reliability shifts from reactive maintenance toward predictable asset performance.

Reaching higher maturity within a single facility, however, introduces a new challenge. The next step is scaling that capability across the organization. This is where many reliability programs begin to stall, as differences between sites, processes, and governance structures can prevent the closed-loop system from operating consistently at enterprise scale. The next section examines why scaling often fails and how organizations can extend reliability capability beyond individual plants to create system-wide impact.

Scaling the Closed Loop Across the Enterprise

Many organizations prove the **Detect > Diagnose > Action > Improve** model in a pilot plant by preventing failures, documenting savings, and building early momentum. However, when expansion begins, performance often plateaus. Results start to vary by site, standards drift, and metrics lose consistency. The system that worked in one location no longer produces uniform impact across the broader network.

Scaling doesn't fail because the closed loop is flawed, but because enterprise alignment is missing.

Challenges to Be Aware of In Scaling and What They Can Do To Your PdM Program

Scaling Challenge	What Happens in Practice	Result
Site-to-site inconsistency	Each facility develops its own thresholds, workflows, and reporting standards. What qualifies as high risk in one plant may be considered routine in another.	Enterprise visibility becomes fragmented and leadership cannot compare risk consistently across sites.
Over-customization	Local optimization overrides standardization. Plants modify workflows and processes to solve short-term problems.	Benchmarking becomes difficult and cross-site learning breaks down. Multiple versions of the closed loop begin operating at once.
Data overload without governance	As monitoring expands, the volume of alerts and data grows rapidly. Without clear prioritization frameworks and ownership, teams struggle to manage the signal.	The signal-to-noise ratio declines and teams become reactive to alerts instead of aligned around risk.
AI deployed without standardization	Analytics models are configured differently by site or applied without shared diagnostic frameworks and asset taxonomies.	Outputs become difficult to compare or trust. Intelligence increases variability instead of reducing it.

Enterprise scaling requires more than deploying technology across additional sites. It requires shared standards, governance, and accountability that ensure the closed-loop system operates consistently across the organization.

The Enterprise Scaling Blueprint

Scaling the closed loop across multiple sites requires structure, not just expansion. Enterprise reliability succeeds when standards, visibility, and accountability are aligned across the network. The following blueprint outlines the core elements that allow reliability systems to scale without losing consistency.

Scaling Element	What It Looks Like in Practice	Why It Matters
Standardized Diagnostic Framework	Shared fault classifications, severity levels, and risk scoring across all facilities.	Ensures alerts mean the same thing everywhere and enables consistent prioritization.
Shared Asset Taxonomy	Common equipment categories, naming conventions, and criticality rankings.	Allows meaningful benchmarking and enterprise-level analytics.
Centralized Reliability Intelligence	Cross-site reliability reviews and dashboards that aggregate data from all plants.	Gives leadership visibility into systemic risk and emerging patterns.
Plant-Level Execution	Individual plants retain ownership of execution, with the flexibility to adapt based on their specific challenges and level of maturity.	While corporate should provide a strong framework for scaling PdM across the organization, each plant operates differently.
Human Oversight and Ownership	Director or VP-level reliability leadership upholds standards. They have consistent reviews with vendors on performance and what can be improved.	Prevents process drift and maintains diagnostic consistency.
Cross-Site Benchmarking	Plants compare reliability performance, intervention timing, and failure trends.	Accelerates learning and replicates best practices across facilities.

When these elements operate together, scaling improves consistency and enterprise control rather than introducing variability.

Metrics at Scale

When reliability expands across multiple sites, measurement must evolve as well. Site-level metrics such as local downtime or alert counts provide limited insight. Leadership needs indicators that show whether risk is being reduced across the entire network.

Enterprise Metric	What It Measures	Why It Matters
Cross-Site Repeat Failure Reduction	Whether the same failure modes continue appearing at different facilities.	Indicates whether systemic defects are being eliminated rather than repeatedly repaired.
Enterprise Avoided Downtime Trend	Total verified downtime prevented across the organization over time.	Shows whether predictive maintenance is compounding financial value.
System-Wide Anomaly Reduction %	Decline in abnormal operating conditions detected across monitored assets.	Reflects decreasing variability and improving operational stability.
Mean Time to Comment (MTTC)	Time from alert creation to first meaningful human acknowledgment.	Indicates adoption and engagement with the system, showing how quickly teams are reviewing and acting on insights.
Mean Time to Repair (MTTR)	Time from alert creation to confirmed resolution or closure.	Reflects operational effectiveness and how efficiently issues are resolved after detection.

Together, these metrics shift the focus from isolated plant performance to enterprise risk control. They reveal whether reliability improvements are compounding across the organization or remaining confined to individual sites.

Case Study: Scaling Reliability at Boise Cascade

Boise Cascade operates a network of manufacturing facilities producing engineered wood products across the United States. As the company began exploring predictive maintenance, leadership faced a challenge common to many organizations: maintenance practices varied significantly between plants, and there was no standardized way to monitor asset health across the enterprise.

To evaluate predictive maintenance, two facilities piloted the platform from AssetWatch. The pilots demonstrated that early fault detection could reduce unplanned downtime and provide actionable insights for maintenance teams. The challenge then shifted from proving value at a single site to scaling the capability across the organization.

The table below illustrates how Boise Cascade progressed through several stages of reliability maturity during this rollout.

Reliability Stage	What Boise Cascade Experienced	What Changed
Stage 1: Reactive Maintenance	Maintenance practices varied by site. Many assets were aging and reactive maintenance was common. Limited condition-based monitoring experience existed across the organization.	Failures often dictated priorities and asset health visibility was limited.
Stage 2: Monitoring Expansion	Two facilities piloted predictive maintenance technology and began comparing insights and results. Sensors provided new visibility into asset conditions.	Teams began detecting faults earlier and documenting initial avoided downtime events.
Stage 3: Operational Integration	Leadership approved enterprise expansion after pilot success. Over 13 months, monitoring scaled from 1 to 17 facilities and from 96 to 1,224 measurement points.	Predictive monitoring became a consistent reliability standard across the organization.
Stage 4: Process & Capability Development	Dedicated Condition Monitoring Engineers supported training, diagnostics, and data interpretation. Plants received site-specific onboarding and asset prioritization guidance.	Maintenance teams learned how to translate predictive insights into maintenance decisions.
Stage 5: Enterprise Alignment	Plants began sharing results and lessons learned across the organization. Success stories and early wins helped drive adoption.	Collaboration increased between sites and predictive maintenance became embedded in daily operations.

Key Enablers of the Transition

Several factors helped Boise Cascade successfully move from pilot testing to enterprise implementation:

- **A clearly defined pilot-to-enterprise adoption strategy.** Corporate established a structured rollout plan, gathering data on each facility to understand plant-level maturity and tailor the approach accordingly. Success was actively monitored and communicated across sites to build momentum, while regular engagement with key stakeholders enabled teams to gather feedback, review alerts, capture financial impact, and **reinforce progress by celebrating wins.**
- **Dedicated reliability expertise** through Condition Monitoring Engineers who supported diagnostics and training.
- **Plant-specific onboarding** that helped each facility identify critical assets and establish monitoring strategies.
- **Cross-site collaboration** that allowed facilities to share lessons learned and early success stories.
- **Gradual scaling** that proved value before expanding the program enterprise-wide.

Over a 13-month period from November 2023 through December 2024, the program expanded to **17 manufacturing sites and more than 1,200 monitored measurement points,** creating a consistent predictive maintenance foundation across the organization.

This transition shows how organizations move from monitoring expansion to operational integration. Early pilot success demonstrated the value of predictive maintenance, but scaling the capability required training, collaboration, and a structured rollout that aligned teams across the enterprise.

The Bottom Line

Closed-loop systems eliminate technical gaps by providing the visibility and diagnostics needed to detect problems early. Maturity eliminates organizational gaps by ensuring insights consistently translate into disciplined action and measurable results.

Zero downtime is not the result of better maintenance alone. It comes from engineered visibility, standardized decision-making, financial validation, cultural ownership, and systems that learn and improve over time.

At that point, reliability no longer operates as a cost center. It functions as an enterprise risk management engine and, ultimately, a source of sustained strategic advantage.

See how AssetWatch helps organizations scale reliability from isolated wins to enterprise-wide operational impact.

Get started today.

Sources:

1. McKinsey Digital. (2015). Industry 4.0: *How to navigate digitization of the manufacturing sector.* https://www.mckinsey.com/~/media/McKinsey/Business%20Functions/Operations/Our%20Insights/Industry%2040%20How%20to%20navigate%20digitization%20of%20the%20manufacturing%20sector/Industry-40-How-to-navigate-digitization-of-the-manufacturing-sector.pdf

2. *Maintenance: ZIPDO Education Reports 2026.* (2026, February 12). ZipDo. https://zipdo.co/maintenance-statistics/

3. Grant, T. (2026, February 13). *Predictive Maintenance Statistics.* Gitnux. https://gitnux.org/predictive-maintenance-statistics/

4. Woollacott, E. (2025, October 29). Manufacturers report millions in losses as downtime wreaks havoc on operations. *IT Pro.* https://www.itpro.com/infrastructure/manufacturers-report-millions-in-losses-as-downtime-wreaks-havoc-on-operations