Effective Real-Time Monitoring of Blast Furnaces With Time Series AI Platform

Suhas Mehta[1], Namrata Rao[1], Dr. Nikunj Mehta[1], Achinta Haldar[2], Richard Zhang[2], John D'Alessio[2]

[1]Falkonry Inc.
20286 Vista Ct., Cupertino, CA 95014, USA
Phone: +1 (408) 461 9286
Email: suhas.mehta@falkonry.com, namrata.rao@falkonry.com, nikunj.mehta@falkonry.com

[2]Stelco Inc.
2330 Regional Road #3, Nanticoke, ON N0A 1L0
Phone: +1-519-587-4541
Email: achinta.haldar@stelco.com, richard.zhang@stelco.com, john.dalessio@stelco.com

ABSTRACT

Effective real-time monitoring is crucial for modern iron and steel operations to prevent issues like cold furnaces or unscheduled shutdowns, which leave signatures in high-volume real-time data. Conventional monitoring often misses these signs.

Artificial intelligence enhances our ability to detect these signatures, helping to avoid lost production, improve quality, and increase campaign duration. Effective AI must offer significant avoidance opportunities, require minimal attention from operations and maintenance, and cover most operational issues to gain trust and advance digital manufacturing. Therefore, measuring and communicating this effectiveness is vital for digitalization in the steel industry.

This paper introduces a new metric for judging the effectiveness of real-time monitoring in digital transformation projects, regardless of the technology. We illustrate this metric with multiple real-world case studies of Time Series AI platforms applied by steelmakers globally.

Keywords: Monitoring, Real-time, Artificial Intelligence, alert prioritization, data-centric maintenance

INTRODUCTION

Effective real-time monitoring is vital for modern iron and steel production. Highly dynamic systems like blast furnaces have thousands of signals containing early indicators of asset degradation and off-normal behavior that precede adverse events. However, conventional monitoring is mostly based on static thresholds and manual data review, causing both excessive alarms and missing failure precursors. This noisy, reactive environment overwhelms operators, hindering timely action. Conventional monitoring is insufficient for the complexity of modern ironmaking, as human operators are simply unable to track all signals or identify subtle, emerging anomalies across intricate steelmaking systems. Therefore, monitoring must be augmented with AI to achieve the necessary level of effectiveness.

AI-based monitoring functions as "a thousand pairs of eyes," capable of effectively monitoring the thousands of signals across a line. To be effective, this AI monitoring must:

- Cover a broad range of operational failure modes, including subtle and systemic anomalies.
- Minimize cognitive load by delivering a manageable number of high-confidence, actionable alerts.
- Integrate into existing operational and maintenance workflows.

The ultimate goal is to generate monitoring results that operators trust, thereby facilitating timely, data-centric interventions. This paper proposes a three-part framework to modernize monitoring: 1) a metric for tracking digital transformation progress, 2) an AI platform to enable rapid, capital-free progress, and 3) a standard operating procedure for leveraging daily AI results.

FOUNDATION

To quantify progress in digital transformation, regardless of the real-time monitoring technology, we introduce the Average Monitoring Reduction Ratio (AMRR). This new metric measures the effectiveness of automated monitoring in watching all data flows continuously while minimizing the cognitive load on engineers to review results.

AMRR specifically quantifies how much automated monitoring would reduce the need for manual signal review. It is defined as the ratio of signals automatically monitored in real-time to the signals that still require active human review before action is taken.

$$AMRR = \frac{Total\ number\ of\ monitored\ signals}{\#\ of\ signals\ requiring\ active\ manual\ review}$$

AMRR is recommended to be calculated over a realistic timescale, such as a shift or a day, as this is appropriate for establishing the denominator of the metric. However, alternative timescales can be used for long-term process improvement analysis.

As a key efficiency KPI, the AMRR can be tracked across assets, lines, or sites and monitored over time, alongside alert quality and action rates. An AMRR significantly above 10 demonstrates that automation effectively replaces continuous manual monitoring with the review of summarized results (anomalies and patterns), indicating a corresponding reduction in operator workload. Higher AMRR values signify greater scalability of the monitoring approach. The ideal value is infinity, which represents the complete elimination of the need for human review. Achieving this optimal state requires the monitoring strategy to be validated for its prescriptive guidance capability and accuracy.

Because it does not depend on a specific use case or machinery, and is not a measure of accuracy (which is frequently challenging to assess due to a lack of labeled data), AMRR is a strong candidate for consistent use across various AI tools and sites and for tracking process evolution over time.

BACKGROUND

Time Series AI platforms are an advanced real-time monitoring technology, which involves four steps:1) data ingestion, 2) analysis, 3) alert generation, and 4) operationalization.

1. Data Ingestion

Iron and steelmaking operations can generate anywhere between 500,000 and 100,000,000 data points per day, depending on system configuration and signal sampling rates. As a real-time monitoring technology, an AI platform must ingest live time series data via MQTT streams and Parquet file transfers, while preserving native resolutions from milliseconds up to hours. Signals can be either numerical (e.g., current, temperature, speed, vibration) or categorical (e.g., valve state, switch position). In order to win the trust of plant operations teams, an AI platform should explain its results in ways that are easily understood, which requires rich and responsive visualization of time series data at multiple resolutions, while highlighting anomalies and patterns.

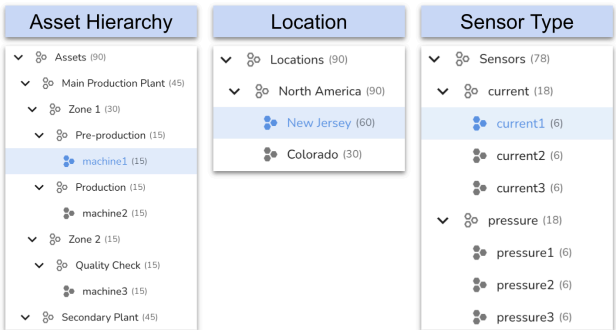

Figure 1: Signal contextualization via flexible, standards-aligned asset hierarchies

When dealing with thousands of signals across lines, it is essential that context is adequately included in visualizations. Standards-based asset hierarchies aligned with ISO 14224 and ISA-95 can be used for this purpose and enriched with contextual metadata such as geographical location, asset criticality, sensor type, equipment type, ownership, and manufacturer (Figure 1). This kind of structuring of signals also enables prioritization of assets for review (root cause analysis) and action (proactive intervention) across complex steelmaking systems, including blast furnaces, reheat furnaces, and finishing mills. The same context is also required for shift handover and daily summary reports with fault hypotheses, allowing operators and engineers to quickly review, validate, and act on likely root causes.

2. Analysis

An AI platform provides a highly automated and resource-efficient means of locating anomalies and patterns in signals. Techniques such as self-supervised deep learning that automatically models all observed behavior and detects anomalies in signals in real time, without requiring any labels [8] are considered to be well-suited to the needs of real-time process and asset monitoring. An AI platform can continuously score each signal for anomalousness as a unitless scalar that can be easily compared across signals. Since these anomaly scores themselves are numerical time series, they can be ingested, visualized, and further analyzed for automated alert generation in the AI platform. Univariate anomalies can be aggregated to create asset-level anomaly views, providing a holistic perspective on potential operational issues and highlighting active conditions. Our prior work [2] categorizes such anomalies into four dominant classes that are highly relevant to iron and steelmaking operations, supporting early warning and condition-based operational decision-making.

The AI platform should also provide tools to identify specific condition patterns in the form of multivariate anomalies. This is highly useful to detect recurring patterns (e.g., strip break, channeling, hanging) and automate root cause analysis from process and asset data [9]. The AI platform can also be trained to detect early warnings of failure patterns, which arise as subtle breakdowns in multiple signals that go undetected by conventional monitoring systems and processes. Such early warnings provide minutes to days for targeted intervention and reduce the risk of unplanned downtime or forced shutdowns. The condition assessment is itself a categorical time series, which, too, can be further processed in the AI platform, e.g., to produce alerts.

Additionally, physics-based models that translate live operational data into domain-specific condition indicators are important to use with AI platforms. These indicators are also time series, which can be further processed in the AI platform, e.g., to produce alerts.

3. Alert Generation

Time Series AI platforms should not only look for patterns and anomalies but also trigger human workflows. In order to translate numerical and categorical time series data into such triggers, it is necessary for the AI platform to process AI outputs through an advanced rules engine that extracts stable signals from probabilistic AI output. This rule engine transforms analog quantities, such as raw signals, anomaly scores, and condition labels, into discrete events [3].

The rules engine acts as an operational layer by bridging the gap between anomaly detection technology and practical operations requirements. Unlike standard CBM systems that react blindly to raw signal amplitude, this engine synthesizes probabilistic AI insights with deterministic physical constraints and operational states. By rigorously filtering transient noise through temporal denoising and enforcing context-aware suppression, it achieves a superior signal-to-noise ratio. This capability elevates rules from a simple notification mechanism into a robust decision-support framework, ensuring that only high-confidence, physically validated conditions are notified to the operator.

3.1. Denoising Alerts: Conventional monitoring systems rely on static thresholds, implicitly assuming that brief excursions beyond limits are both abnormal and actionable. In real steelmaking operations, equipment is routinely operated near or beyond nominal design envelopes, causing threshold-only rules to generate high-frequency, low-value alarms. The result is alert overload, operator desensitization, and loss of trust in the monitoring system. It is therefore necessary to process anomaly scores, pattern IDs and domain-specific indicators through an advanced rules engine by introducing four orthogonal rule parameters that control *when* a condition is considered meaningful and *how often* it is allowed to interrupt human attention (Figure 2). Together, these parameters act as a temporal, spatial, and frequency-based filter that suppresses transient and localized noise while preserving persistent, asset-relevant behavior.

1. **Threshold/Condition** - Defines the basic abnormality criterion: a numeric limit for numerical signals or a discrete state for categorical signals. This establishes what constitutes alertable behavior but does not, by itself, generate alerts.
2. **Persistence (Density)** - Specifies the minimum percentage of time within an evaluation window that the threshold or condition must be satisfied. This parameter removes short-lived spikes, sensor chatter, and momentary process transients that are not operationally meaningful.
3. **Spread (Coverage Depth)** - Defines how many input signals must simultaneously meet the threshold or condition. This spatial filter suppresses single-sensor artifacts and elevates conditions that are correlated across multiple components, subsystems, or measurements.
4. **Alert Frequency** - Controls how often alerts are emitted once all criteria are satisfied, preventing repeated notifications for the same underlying condition and bounding alert volume to human-manageable levels.

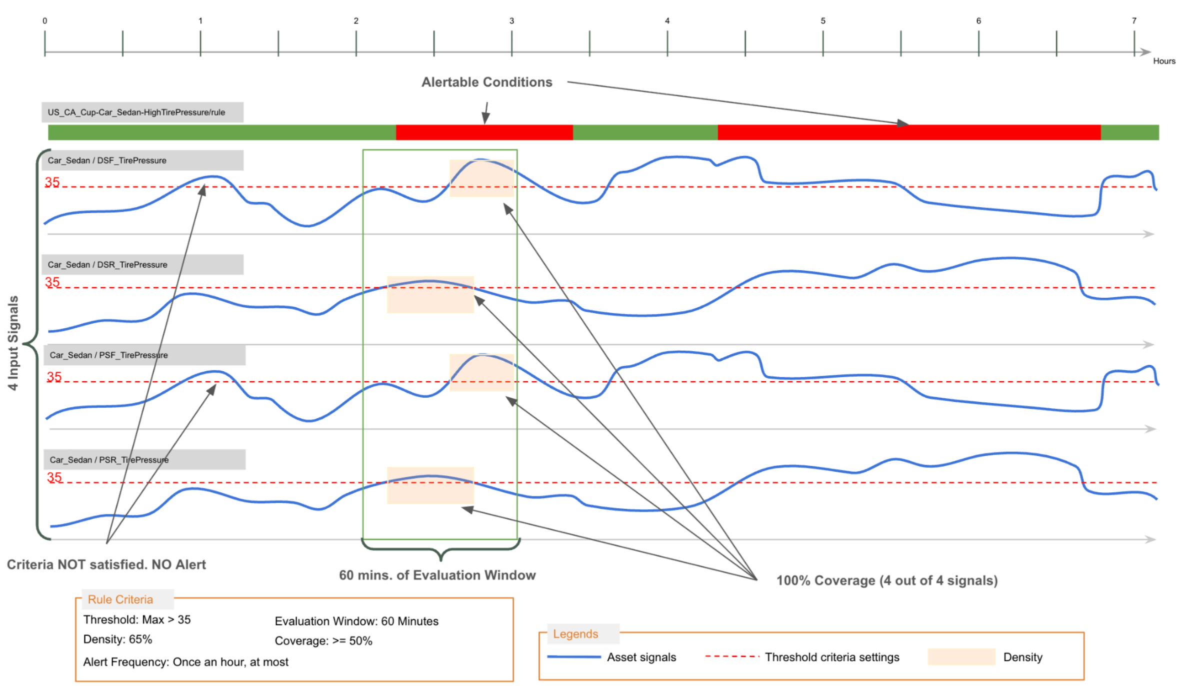

Figure 2: Illustration of how rule parameters are evaluated together to convert raw signal behavior into denoised, actionable alerts

The AI platform should support both simple and compound (chained) rules. Each component rule applies the same core denoising parameters but at different semantic levels, enabling progressively higher-confidence alerts. Rules can be applied to both numerical and categorical time series in the form of thresholds on its statistical properties

Advanced Rules, i.e., those with denoising parameters, do not merely identify violations; they encode persistence, correlation, and alerting discipline into the evaluation itself to make for an alerting environment that is ISA-18.2 compliant, which limits human alarms to a maximum of 150 alarms/day (300 max.) per operator position [5], emphasizing that "the key steps of removing things that are not really alarms, training the operators to respond, and monitoring system performance provide most of the benefits " [4]. Advanced rules apply these human-factors principles by controlling alert persistence, signal coverage or spread, and firing frequency, limiting alerts based on asset criticality.

3.2. Temporal Denoising: When an advanced rule evaluates one numeric or categorical signal against defined criteria, it uses persistence to suppress transients. Figure 3 illustrates a numerical rule in which an alert is generated only when Motor_Current2 exceeds 250 Amps for at least 50% of a 5-minute evaluation window, suppressing transient current spikes and triggering alerts only for sustained overload conditions.

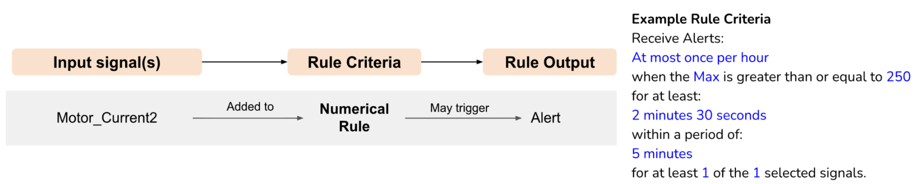

Figure 3: Illustration of a single-signal numerical rule

3.3. Spatial Denoising: Complementing the temporal filtering mechanisms, advanced rules also incorporate spatial denoising to validate the scope of a detected condition. This logic mandates that deviations occur concurrently across multiple distinct signals to confirm an event, thereby distinguishing systemic process shifts from isolated sensor artifacts. Anomaly rules evaluate AI-generated normalized anomaly scores using the same persistence and spread logic applied to raw signals. This prevents isolated or low-confidence AI detections from generating alerts. Figure 4 illustrates a multi-signal anomaly rule that raises an alert only when at least 2 of 3 anomaly score values exceed 6 for 20% of a 25-minute evaluation window, converting noisy anomaly streams into stable, interpretable, and actionable events.

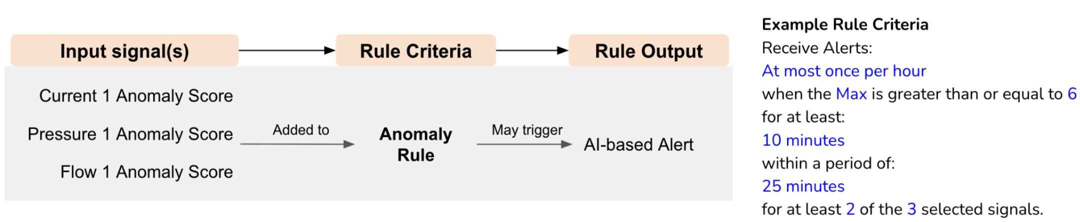

Figure 4: Illustration of a multi-signal (numerical) anomaly rule

Figure 5 illustrates a multi-signal condition rule in which an alert is generated only if any of the 23 AI condition models detect a *Warning* state for at least 16% of a 30-minute evaluation window, filtering short-lived condition detections and surfacing only persistent, operationally relevant events.

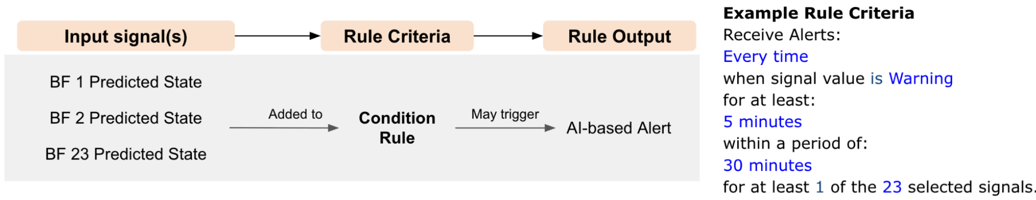

Figure 5: Illustration of a multi-signal (categorical) condition rule

3.4. Contextual Alert Suppression: A compound (chained) rule is a specialized form of an advanced rule in which inputs are categorical rule outputs (true|false) themselves (Figure 6). Rather than operating directly on raw signals, compound rules evaluate the results of two or more upstream rules, enabling logical gating and contextual suppression of alerts.

Compound rules generate alerts only when multiple prerequisite conditions are simultaneously satisfied, ensuring that alerts represent meaningful, multi-stage events rather than isolated or expected behaviors. This structure is particularly effective at suppressing noise caused by maintenance activities, planned startups and shutdowns, transient operating modes, or sensor downtime.

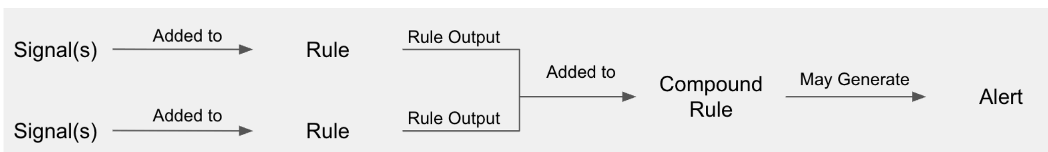

Figure 6: Illustration of a compound rule

A common use case is triggering an alert only when an AI-detected anomaly is present, and the asset is confirmed to be in a valid operating state. By explicitly encoding operational context, compound rules prevent spurious alerts that would otherwise be generated during non-actionable periods.

APPROACH

Rules are effective only when its outputs are consistently translated into operational action. First order of business is criticality determination using practical engineering criteria such as redundancy and consequence of failure. Redundant systems (e.g., circulation fans) are of lower criticality, while single-point-of-failure systems, like the blast furnace hearth, are the highest due to catastrophic shutdown risk. This can be encoded in a Criticality Tree, mapping assets to discrete levels (Urgent, High, Medium, Low) and integrating directly into rules and alert logic so that alerts from higher-criticality assets are automatically elevated.

1. Critical (Urgent) Alerts

Urgent alerts require immediate action and should be delivered to the operator in real time using, say webhook or an MQTT topic, while other alerts are non-critical and therefore aggregated into shift handover or daily summaries to reduce alert fatigue. Critical alerts are automatically escalated into high-priority Work Orders (WO) within the OT or CMMS environment. Examples of critical alerts include, but are not limited to,

- **Infrastructure alerts** (e.g., gateway, network, or power failures) are treated as urgent because they disable monitoring and must be resolved promptly to avoid blind operation
- **Blast furnace channeling alerts:** Early detection of channeling conditions enables operators to intervene before burden instability escalates into hanging, slips, or productivity loss, reducing the risk of severe furnace disturbances and unplanned downtime

2. Non-Critical Alerts
Non-critical alerts are delivered via asset-centric summaries (shift handover or daily reports), connecting detection, prioritization, and execution. Instead of individual signal alarms, the hybrid monitoring solution groups alerts by asset and presents a prioritized Asset-Alert list aligned with maintenance and operations workflows.

2.1. Asset Prioritization: SMEs define a subset of high-impact assets or sub-assets for active monitoring, focusing on equipment with the greatest safety, production, or quality consequences. Less critical assets can be phased in as data quality and organizational readiness improve, ensuring early deployments deliver value without overwhelming operations.

2.2. Recurrence-Based Prioritization: Alert recurrence within a rolling window (e.g., 24 hours) is used to distinguish chronic degradation from transient noise. Figure 7 shows that the assets with high alert counts or severe conditions are elevated as **Action items** for immediate attention, while those with lower counts or emerging patterns are designated **Watch items** for ongoing monitoring and trend observation. Repeated or persistent alerts are escalated, while isolated, self-resolving events remain lower priority. This prevents overreaction to momentary disturbances while ensuring sustained issues receive timely attention.

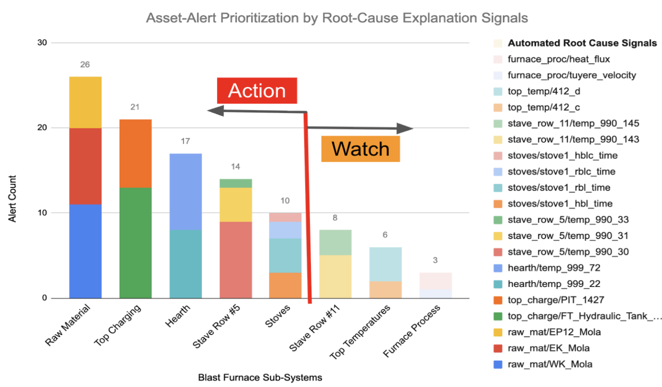

Figure 7: Illustration of Action and Watch lists derived from daily or shift-based alert summaries

3. Shift Handover Report or Daily Alert Summary
Non-critical alerts are best reviewed on a per-shift or per-day basis rather than acted on immediately. In blast furnace operations, for example, alerts are typically reviewed once daily at the start of a shift or production day to support predictive maintenance planning. By consolidating alerts into structured summaries, maintenance and operations teams can focus on the highest-impact interventions, allocate resources efficiently, and maximize uptime, ensuring AI-driven detection is consistently translated into actionable decisions. A typical alert summary report includes:

- A prioritized view of assets with action and watch items, (Figure 7)
- A ordered list of assets clearly distinguishing AI-based alerts from logic-based alerts (Figure 8)
- Key contributing signals for each alert, enabling automated root-cause insight (Figure 7)
- Hourly trends of critical signal statistics (e.g., mean, standard deviation) to highlight assets or subsystems requiring attention (Figure 9)
- A 45-day alert trends to surface recurring or escalating conditions and support longer-term reliability and maintenance planning (Figure 10)

Asset	AI Alerts	Threshold Alerts	Total Alerts	Contributing Signals
Raw Material	28	61	89	- WLH_Coke_Disc_Time o SD: 5.965 | Mean: 114.799 | Count: 1 - EP34_Mola o SD: 0.843 | Mean: 0.576 | Count: 3 - WK_Mola o SD: 2.118 | Mean: 9.027 | Count: 34 - EK_Mola o SD: 1.112 | Mean: 6.517 | Count: 14 - ELH_Coke_Disc_Time o SD: 13.988 | Mean: 127.469 | Count: 37
Top Parameters	3	0	3	- TE_412_A o SD: 23.628 | Mean: 107.8 | Count: 3

Figure 8: Illustration of how a single asset may have a combination of AI and threshold alert

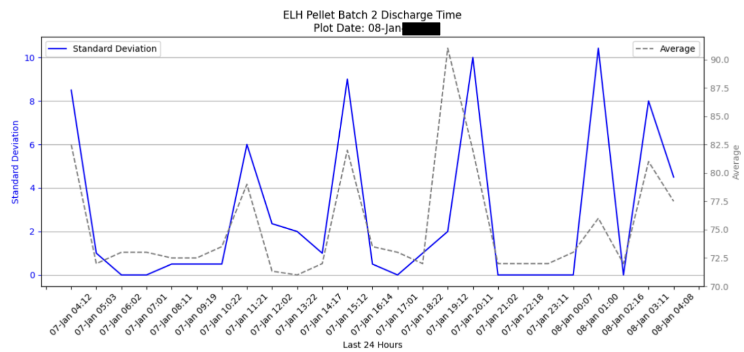

Figure 9: Example of hourly trends on a root-cause explanation signal

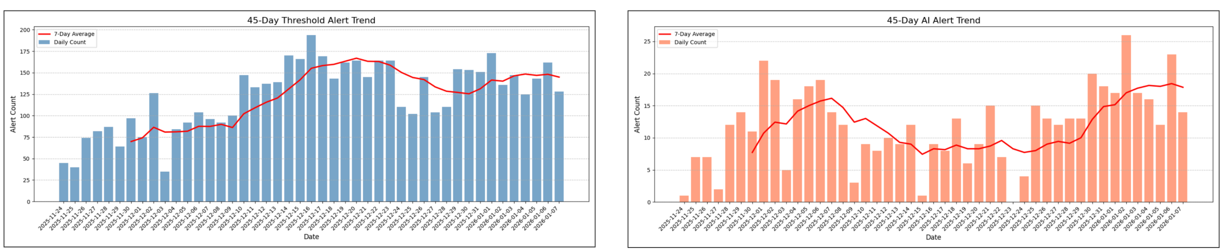

Figure 10: 45-day trend of threshold-based and AI-based alerts

4. Digitalization Workflow Adaptation

Figure 11 shows a traditional pre-digital workflow where raw SCADA, PLC, historian, and enterprise signals flow directly to a process engineer. Large volumes of unprioritized blast furnace data must be manually reviewed alongside ambient conditions and product recipes. Analysis is intermittent, with issues often identified only during twice-weekly, 60-minute stand-ups, relying on trends and anecdotal observations. This workflow is time-consuming, reactive, and dependent on individual expertise rather than systematic prioritization.

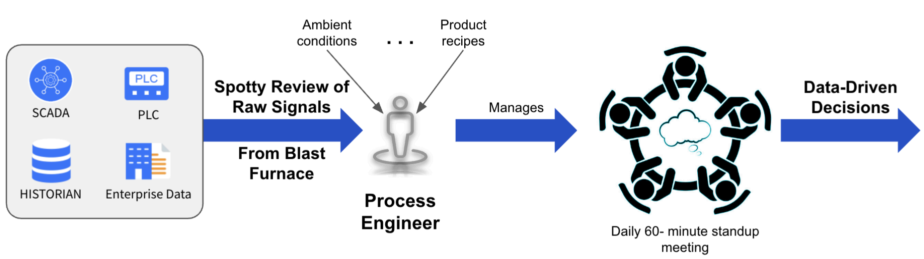

Figure 11: Traditional operations workflow for managing a blast furnace

Figure 12 illustrates the transformed workflow: raw signals from SCADA, PLCs, and historians stream into the AI platform for continuous analysis, generating prioritized alerts and daily summaries. This shifts the process engineer's role from manual data aggregation to the review of curated, asset-aware insights. Consequently, lengthy bi-weekly meetings are replaced by sub-15-minute daily sessions focused on high-impact, data-centric interventions. This evolution establishes a proactive monitoring culture, prioritizing early intervention over retrospective discovery.

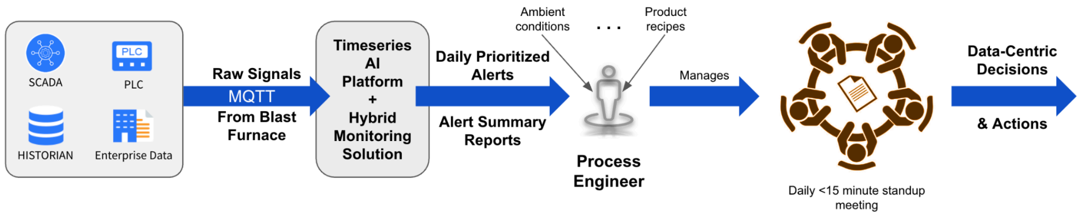

Figure 12: Digitally transformed operations workflow for effective blast furnace monitoring and management

EVALUATION

We examine three distinct applications: dual-layer process monitoring on a North American Blast Furnace, high-frequency vibration analysis on Reheat Furnace fans, and automated strip break classification. We present an evaluation of AMRR across these use cases to assess the effectiveness of automated monitoring.

1. Blast Furnace Monitoring

A leading North American steelmaker deployed a Time Series AI platform to monitor blast furnace operations on more than 229 prioritized PI historian signals. These signals spanned critical blast furnace subsystems, including burden charging, gas flow, stoves, staves, furnace process, and material handling, enabling continuous, asset-centric visibility into both localized equipment behavior and overall furnace dynamics.

The first stage of automated monitoring employed static thresholds to detect anomalies across both slow- and fast-sampled signals. This deterministic approach provided easily understood alerts aligned with established operating limits and helped operators validate the reduction in alert load as well as build confidence in the system while addressing known, critical failure modes. The second stage of monitoring was to activate AI-driven anomaly detection, intended to identify subtle deviations that precede furnace instability and malfunction. Using rules as explained in the previous section, AI outputs were translated into operationally relevant alerts.

Initially, more than 1,100 blast furnace signals were streamed from the PI historian into the AI platform over MQTT. Through SME-driven prioritization and early system learning, this set was narrowed to 229 high-impact signals spanning critical blast furnace subsystems. These priority signals were then continuously monitored to produce approximately three to five signals of interest daily across 1 or two alerts each. Manual review effort was narrowly focused on these signals of interest and required minimal effort or cognitive load. The resulting daily AMRR of 45.8 (229 automated signals / 5 manually reviewed signals) has proven to be the difference between completing a daily operational review in 15 minutes vs. not being able to conduct regular operational reviews at all.

Table 2: Distribution of independent adverse events

	AI Anomaly Events	Threshold Anomaly Events
Process upsets (Cold Furnace Operation)	2	2
Breakdown / Unscheduled Shutdown	0	7
Slowdown	3	9

Total	**5**	**18**

Over seven months, by following this daily process, this steelmaker was able to take timely action to avoid 4 cold furnace events, 7 breakdowns or unscheduled shutdowns, and 12 production slowdowns, i.e., 23 early warnings of potentially severe adverse events in total. These insights correspond to an estimated 150+ hours of potentially avoided unscheduled downtime or slowdown. The impact is a result of a well-calibrated AI platform and a data-centric operational management process and illustrates the potential for impact in other parts of the steel value chain.

As monitoring maturity has increased, the steelmaker has further expanded to 397 prioritized signals and achieved a daily AMRR of 16.54. This reduction in the AMRR is attributed to the blast furnace instability due to prevailing harsh operating condition (severe winter). The steelmaker expects to further improve the AMRR and bring it closer to the previous level by dialing in the right level of sensitivity. This goes to show that AI-based monitoring can be scaled without a corresponding increase in human attention, reinforcing operator trust and enabling durable adoption of AI-driven blast furnace monitoring.

Beyond direct downtime prevention, the system also enabled earlier investigation of emerging process issues, reduced reliance on manual trend analysis, and supported more proactive, data-centric blast furnace operations. Daily alert summary reviews have become an integral part of the operating rhythm, supporting focused 15-minute daily stand-up meetings among process engineers, operators, and maintenance engineers to prioritize actions and track follow-ups.

1.1. Detections and Alerts

Rules using persistence and coverage depth aggregated multiple alerts from related tags (for example, several valve‑timing signals) into single events, simplifying operator review and enabling cross‑validation between AI and thresholds - for instance, a WLH disc‑time threshold corroborating an AI‑detected anomaly in burden distribution timing.

1.1.1. Early Anomaly Detections Compared to Static-Threshold

AI-based monitoring detected long-running anomalies in key blast furnace signals, including the Moisture On Line Analyzer (MOLA), stave temperatures, and tank pressure, several hours before conventional static thresholds would have been breached. In addition, the system identified systemic patterns in coke moisture and pellet discharge times, which were indicative of abnormal furnace movements and uneven burden distribution. These early insights enabled targeted mechanical and IT interventions, providing operators and engineers with actionable information well ahead of potential production disruptions.

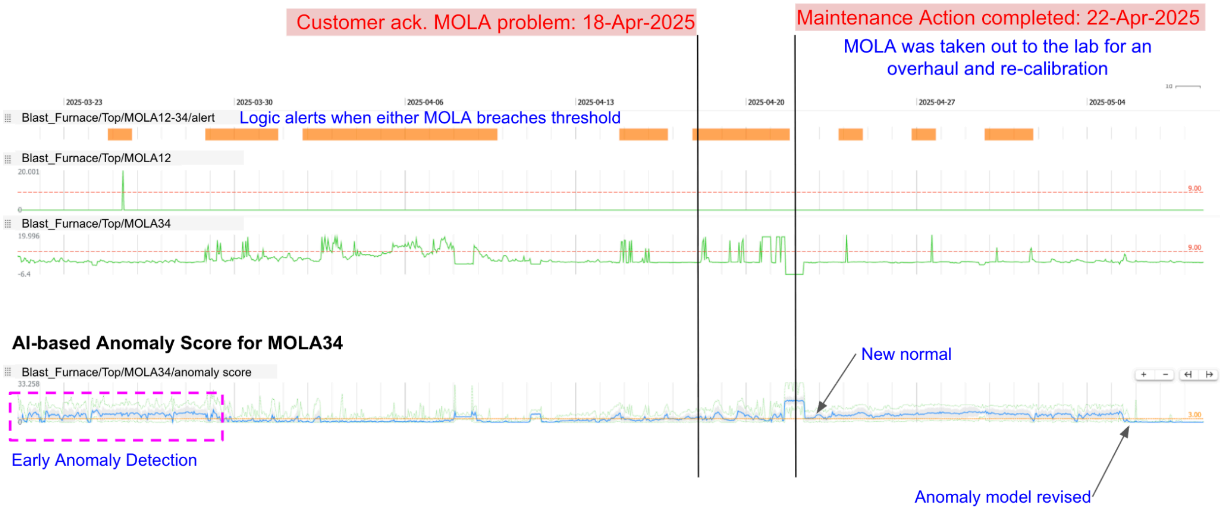

Figure 13: Early insights from MOLA34 anomalies that helped avoid a cold furnace event

Figure 13 shows that the logic-based threshold alerts trigger noticeably later than the high anomalies detected by the AI anomaly model, demonstrating that AI provides earlier visibility into emerging issues. This earlier detection not only avoided a cold furnace event but also built confidence within the blast furnace operations team to investigate high-severity anomaly insights proactively, enabling timely intervention and helping avoid an asset breakdown.

1.1.2. High-severity anomaly when peeling results in low stave temperature

The heatmap in Figure 14 shows AI-detected high-severity anomalies (yellow) in stave temperatures that persist for several hours [2], well before traditional thresholds are breached. Operators observed significant peeling during this period and took preventive action based on the alert, successfully avoiding a process upset like a cold furnace shutdown.

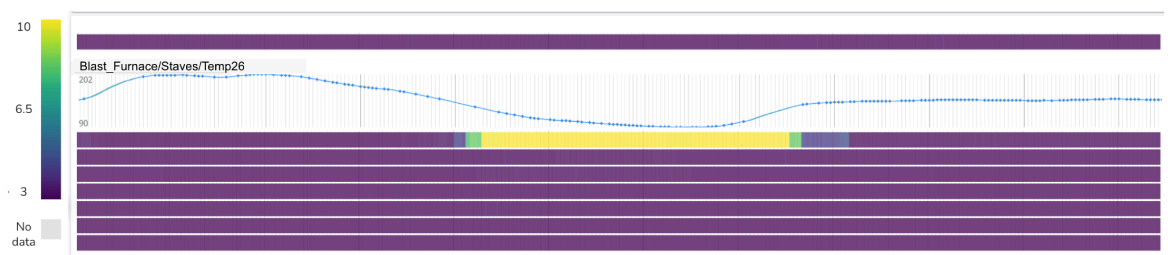

Figure 14: A long-running anomalous stave temperature that avoided a cold furnace event

1.1.3. Short burst increase in the west lock hopper closing time

Figure 15 shows the system detecting a short-burst anomaly [2] in the west-side Lock Hopper closing-time signal, where the valve took longer to close, over an hour. The rising trend and associated high anomaly scores prompted operators to investigate, leading them to zero the tare weight and address the underlying cause, thereby avoiding progression toward a cold furnace event.

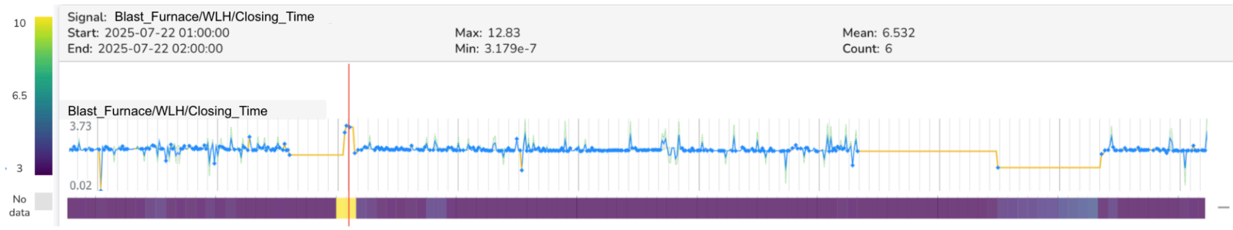

Figure 15: A short-burst anomaly captured by the system that avoided a cold furnace event

1.2. Multivariate Pattern Recognition

Additionally, a multivariate pattern-recognition model was also developed using only a small number of labeled examples to detect both channeling and hanging behaviors in the blast furnace.

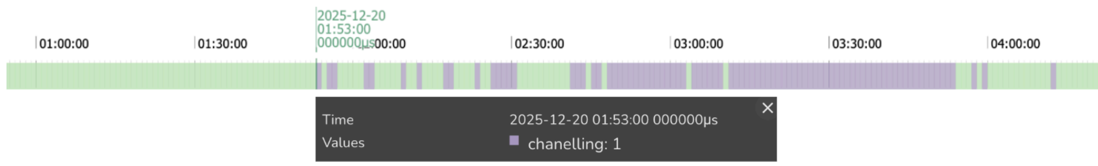

Figure 16: A multivariate pattern recognition model to detect channeling in the Blast Furnace

Alerts on these automated pattern detections in Figure 16 can be fed directly into the operator dashboard, providing a continuously updated, real-time indication of furnace status and unstable burden behavior. This enables operators to recognize and respond to emerging channeling or hanging conditions far earlier than is possible with manual trend review alone, significantly reducing the risk of missed or delayed interventions.

Multivariate pattern recognition enables earlier and more reliable detection of channeling and hanging conditions than manual trend review, significantly expanding the response window for corrective action. By continuously synthesizing correlated signals into a single condition indicator, it reduces missed or delayed interventions, improves consistency across shifts and operator experience levels, and lowers the risk of severe furnace events.

Over time, this capability is expected to reduce cold furnace incidents, slowdowns, and refractory damage, supporting longer campaign life and improved furnace stability. As operators gain confidence in AI-driven early warnings, monitoring efficiency and trust in automated systems increase, laying the foundation for more proactive blast furnace operations.

2. Reheat Furnace and Vibration Monitoring

A European steelmaker deployed a Time Series AI platform on a reheat furnace for monitoring 335 signals covering critical subsystems such as circulation fans and drive systems to detect early anomalies in motor torque, current, vibration, speed, and tension [2].

Real-time monitoring using Time Series AI reduced data to a daily average of 5 prioritized signals, achieving a 67 AMRR. The operational value is clear in the forensic analysis of three September 2023 failure modes, which caused a cumulative production loss of $USD 2.7 million.

- Mechanical Failure: During the same period, a separate mechanical failure involving the "Kulskruvar" (ball screw/feed mechanism) caused 50 hours of downtime, representing a financial impact of USD 1.25 million. The AI platform identified distinct anomaly signatures 13 hours in advance, validating its ability to detect mechanical degradation in non-vibration signals that are often missed by standard condition monitoring.
- Hydraulic System Failure: A catastrophic hydraulic failure resulted in 47 hours of downtime, costing the facility an estimated USD 1.17 million. In the absence of AI guidance, the manual root cause analysis was

labor-intensive and inaccurate. Maintenance teams spent over 75 man-hours diagnosing the issue and erroneously replaced two healthy components before identifying the actual fault. Post-incident analysis revealed that the AI platform had detected multivariate precursors 36 hours prior to failure, offering a window that would have allowed for a planned intervention rather than a reactive shutdown.

- Process Slowdowns: Beyond catastrophic stops, the system demonstrated sensitivity to "soft failures." It detected a "Ström" (flow/current) deviation where signal values dropped to zero, predicting a process slowdown 12 hours in advance. While less severe than a full stoppage, this event still caused a USD 150,000 loss.

Table 3: Avoidable downtime & slowdowns detected by Time Series AI

#	Detection	Downtime & slowdown experienced	Early warning
1	Kulskruvar Downtime	50 Hours	13 Hours
2	Hydraulik Failure	47 Hours	36 Hours
3	Ström Slowdown	9 Hours	12 Hours
	Total	**106 hours**	

The aggregate data from this deployment confirms that a high AMRR does not merely reduce the cognitive load on operators; it directly correlates with the prevention of capital loss. By transforming urgent operational alerts into actionable lead time, the facility could have converted over 100 hours of unplanned downtime into scheduled maintenance tasks.

3. Cold Mill Strip and Weld Break Classification

A third case study at a global steelmaker's Cold Mill applied Time Series AI to a different challenge: the instantaneous classification of strip and weld breaks. These events caused 3–4% (10–15 days) of annual lost production, representing up to USD 3.5M in financial losses per year, and required lengthy manual diagnosis (~400 hours/year).

In the cold mill, frequent weld and strip breaks cause long stoppages in order to fully recover the line and restart production. Because it is very difficult to know the cause of a break, it is not easy to shorten the recovery time. A Time Series AI platform was selected to develop a break classification system to automatically classify the cause based on PLC tags from the cold mill stands.

The solution leveraged self-supervised AI and multivariate classification to:

- Data & Scope: Reduce the input data from 1,350 plant tags to 7 significant electromechanical signals (e.g., Current, Tension, Torque) per stand that reliably predicted break causes.
- Result: Provide an instantaneous, multi-class classification (e.g., Edge Defect, Bad Weld) immediately following the break.
- Impact: This automation transformed a reactive, manual diagnosis process that took hours/days into an instantaneous alert, directly translating to accelerated recovery and a 1% production improvement, valued at USD 1M per year.

Because automated classification surpassed human accuracy, it was possible to avoid any data review to assess the root cause of a break. In effect, almost all the breaks happened on two stands out of the five in the mill, and for each event the AMRR was $34/0 = \infty$.. This application demonstrates the use of Time Series AI to reduce diagnostic load rather than purely monitoring load, achieving high classification quality with a small, relevant subset of data, and providing a significant financial return by minimizing downtime.

CONCLUSION

This paper has presented a three-pronged framework to modernize real-time monitoring consisting of: 1) a normalized metric of digital transformation progress and impact, 2) an AI platform to achieve rapid progress without capital expenditure, and 3) a standard operating procedure that leverages AI results on a daily basis.

We show that AI enables effective real-time monitoring across blast furnaces, reheat furnaces, and cold mills. A Time Series AI Platform extracts and alerts on anomalies and patterns that are too difficult for experts to see, or require offline review or are completely missed by conventional tools. The Average Monitoring Reduction Ratio (AMRR) provides teams a normalized KPI to assess the suitability of AI to their operations and the readiness to scale its application. The metric is a good candidate for use across tools and sites as well as to track progress over time as it is not dependent on a use case or machinery nor is it a measure of accuracy, which is difficult to measure given the lack of labeled data.

In order to achieve the highest possible AMRR, it is necessary to combine different AI methods including logic, self-supervised and semi-supervised learning. The AI platform makes it easy to match problems to these techniques in an operational environment, which can identify specific well-known events as well as surface previously unknown behaviors. Especially, as this AI approach leverages existing PLC and historian data, it eliminates the need to impact production activities or perform significant capital expenditure. This is the ultimate objective of Industry 4.0 to improve production without the need to make any hardware changes to the plant and steelmakers can indeed achieve results in their own operations with technology currently available.

This paper also explains how to manage the operationalization of AI platforms into a standard operating procedure by use of alerting rules. The resulting approach prioritizes the actions of maintenance and process engineering teams as well as avoids cognitive overload. By being able to conduct daily review of operations guided by real-time monitoring results from the previous day, an operations team can significantly improve its hour by hour understanding of the plant, which makes the team more agile and responsive to environmental and market demands.

We demonstrate the versatility of a Time Series AI platform across the full steelmaking value chain, from primary ironmaking to finishing. Real-world results at three different steelmakers across the world validates the three pronged real-time monitoring framework that leverages Time Series AI This framework can move steelmaking from reactive troubleshooting to a proactive, data-centric industry, significantly enhancing equipment reliability and reducing downtime.

REFERENCES

1. CESMII. Enhancing Productivity of Steel Cold Rolling Through Time Series AI. Case Study, 2022. Available at: Improving Productivity of Steel Cold Rolling with Time Series AI. Accessed January 8, 2026.
2. Falkonry, Inc. Prioritized Anomaly Management in Steel Production Utilizing Self-Supervised AI. AISTech 2024.
3. Falkonry, Inc. Definition and Application of Falkonry Rules. Falkonry Time Series AI User Guide. Available at: https://docs.falkonry.ai/userGuide/introTSI/#2-rules. Accessed January 8, 2026.
4. Dunn, D. G., and Sands, N. P., P.E., CAP. "Continuous & Batch Processing: How to Achieve an Effective and Efficient Alarm Management Program." InTech Magazine, International Society of Automation (ISA), March-April 2020.
5. International Society of Automation (ISA). ANSI/ISA-18.2-2016: Management of Alarm Systems for the Process Industries. Research Triangle Park, NC: ISA, 2016. Accessed January 8, 2026.
6. Tennant, M. "Implementing Alarm Management per the ANSI/ISA-18.2 Standard." Control Engineering, 2013. Available at: https://www.yokogawa.com/us/library/resources/media-publications/implementing-alarm-management-per-the-ansi-isa-182-standard-control-engineering/. Accessed January 8, 2026.

7. Emergent Mind. Definition of Real-Time Monitoring. Available at: https://www.emergentmind.com/topics/real-time-monitoring. Accessed January 8, 2026.
8. Singh, C. et al. Automated, fast, multi-timescale Time Series AI, Advanced Process Control – Smart Manufacturing Conference, 2022
9. Waters, C , et al. Transforming metal production by maximizing revenue generation with Operational AI, AISTech 2021