

About the Facility

Location

United States Midwest

Product Capacity

2.5 million tons per year (clinker),
2.7 million tons per year (cement)

Deployment Scope

47 AI models, 2,500+ data tags

First Moth Impact:

~\$500,000 in avoided issues

Alert Response Rate:

>90%

Projected Annual Savings:

\$ 4-6 million



Leading US Cement Producer Saves \$500K in Avoided Issues Caught in the first 4 weeks after deploying UptimeAI Reasoning Engine

The Challenge: When Equipment-Level Monitoring Misses the Root Cause

A leading United States cement producer supplies over 12% of total US cement. About a fourth of this production capacity comes from a major Midwest facility producing 2.5 million tons of clinker and 2.7 million tons of cement annually. As one of the company's largest facilities, the plant's continued success depends on reliable, consistent, and predictable operations.

Plants have traditionally addressed this limitation by leveraging human experts to investigate root causes. Process engineers, maintenance teams, and operations personnel would collaborate to connect the dots between disparate systems. However, as expertise wanes in organizations, this reliance is becoming unsustainable.

The cement producer needed a solution that could identify root causes, not just symptoms, and scale expert-level analysis across all operating areas.

The Solution: AI Reasoning Agents That Think System-Level

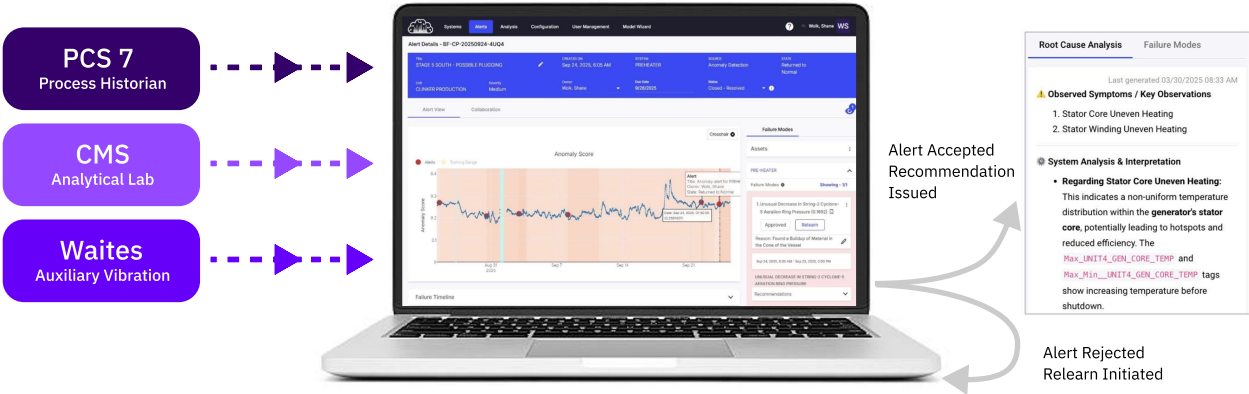
Rather than relying solely on human experts to connect the dots between disparate systems, the cement producer partnered with UptimeAI to deploy AI reasoning agents that analyze operations the way their expert team does—but continuously, across all units, 24/7.

How It Works

The system integrates data from multiple sources to get a complete operational picture:

- Process Control System (PCS 7) and Process Historian
- Computerized Maintenance Management System (CMMS)
- Analytical Laboratory data
- Auxiliary Vibration monitoring
- Temperature monitoring systems

UptimeAI connects sensor, lab, and auxiliary vibration & temp data to get the full system-level view.



Instead of monitoring individual equipment in isolation, the platform analyzes sensor data, lab data, and auxiliary vibration monitoring systems together—taking system-level views to reason cross-functionally like experts.

When the AI detects data quality issues, integrity operating window deviations, or anomalies, it issues alerts along with specific recommendations for how to troubleshoot or mitigate the problem. The model then self-learns through the acceptance or rejection of alerts by operations teams and subject matter experts, continually improving alert quality over time.

Deployment Scope

- The facility deployed 47
- AI models monitoring 2,500+ data tags continuously across multiple production units.

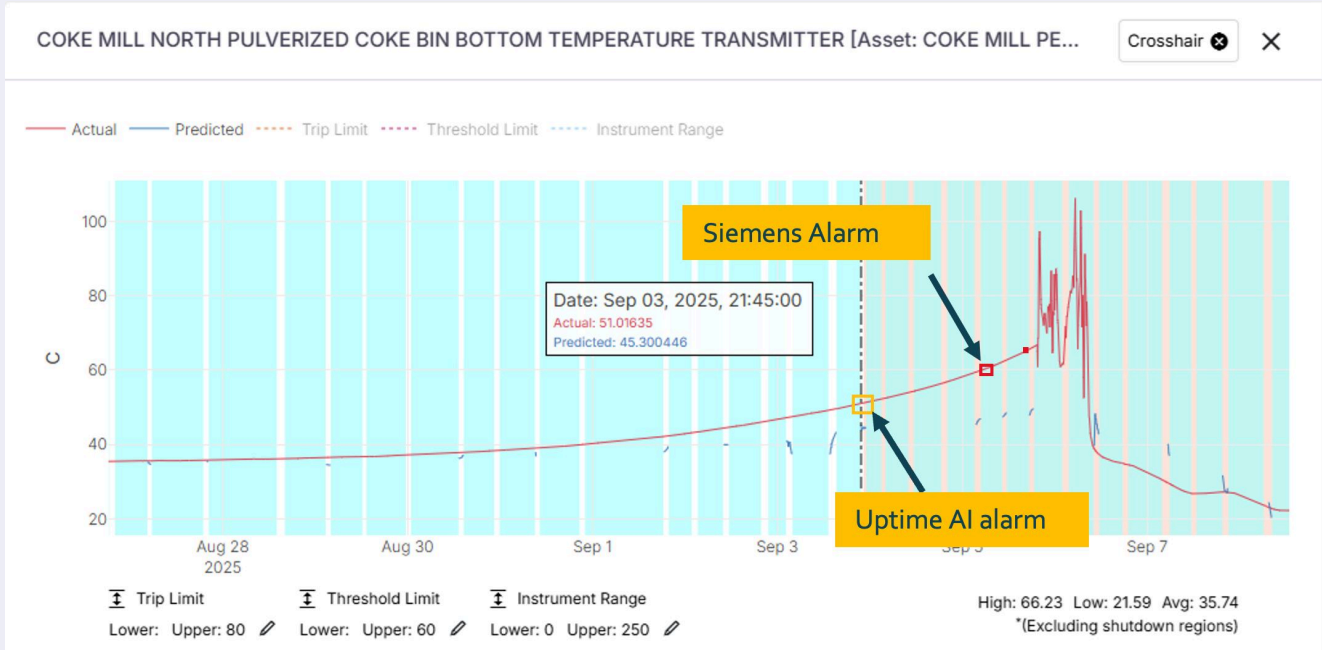
The Breakthrough: Catching Problems Days Before Traditional Systems

The value became clear in a few short weeks of deployment.

An UptimeAI alert was issued for a temperature deviation indicating coke bin smoldering. The critical difference: the alert came days before the control system alarm would have triggered.

The early warning enabled the operations team to take corrective action immediately, avoiding approximately \$75K in estimated damage to the Schenk Coriolis system and a 2-day kiln outage for replacement.

This incident demonstrated what system-level analysis can catch that traditional equipment-level monitoring misses: subtle patterns and correlations across multiple data sources that indicate developing problems well before they manifest as equipment failures.



First Month Results: \$500K in Avoided Issues

In just the first 4 weeks of deployment, the facility prevented multiple significant incidents:

Coke Bin Fire Prevention

- Temperature deviation detected days in advance
- Avoided Coriolis system replacement: ~\$75K

Plugging Event Detection

- Issue identified with sufficient lead time
- Avoided 4-hour kiln line shutdown: ~\$180K

Additional Catches

- Multiple other issues identified and addressed early
- System demonstrated consistent accuracy across various failure modes

Total First Month Impact: ~\$500K

The accuracy and breadth of these catches in just 4 weeks demonstrates significant potential as the deployment expands across the full site.

Why It's Working: High Accuracy Drives High Adoption

The facility achieved a greater than 90% alert response rate from operations teams—a critical indicator that the alerts are trusted and actionable.

This high adoption rate is the result of maintaining low false positive rates through the continuous learning feedback loop. When operations teams accept or reject alerts, the AI learns and improves, keeping precision high and ensuring teams remain engaged with the system over time.

The system has also reduced reliance on individual experts to connect the dots between systems. Expert-level analysis capabilities now scale across all shifts and operating conditions, rather than depending on whether the right person is available when an issue emerges.

Projected Impact: \$4-6 Million Annually

Based on the first 4 weeks of deployment results, the facility projects annual cost savings of \$4-6 million from early fault detection, reduced unplanned downtime, and improved maintenance planning.

This projection is grounded in the demonstrated ability to:

- Detect issues days or weeks before traditional monitoring
- Identify root causes rather than just symptoms
- Prevent cascading failures that lead to extended shutdowns
- Optimize maintenance timing and resource allocation

What's Next: Expanding Beyond Reliability

Current Focus: Mechanical and Production Reliability

The initial deployment focused on mechanical and production reliability challenges, delivering immediate measurable value and building confidence in the approach.

Future Applications: Emissions and Quality Optimization

Emissions Reduction

- Applying system-level analysis to emissions performance
- Supporting environmental compliance and sustainability objectives
- Identifying process adjustments that reduce environmental impact

Product Quality Enhancement

- Optimizing process conditions for consistent product quality
- Reducing variability in final products
- Connecting process parameters to quality outcomes

These applications will leverage the same platform technology and approach that proved successful for reliability, extending the value across additional operational priorities.

The Difference: From Equipment-Level to System-Level

Traditional predictive maintenance monitors individual equipment and flags when parameters exceed thresholds. This approach catches failures but doesn't explain why they occur or how to prevent them from recurring.

UptimeAI's AI reasoning agents analyze data across process control, equipment health, laboratory results, and operational context—reasoning across functional silos the way expert teams do. This system-level approach identifies underlying causes by analyzing patterns and relationships across multiple data sources, then provides specific guidance on how to address problems at their root.

The continuous learning component ensures the system becomes increasingly accurate and relevant to the specific operational characteristics of each facility over time.

About UptimeAI

UptimeAI provides AI reasoning agents that help industrial operations teams make goal-oriented, real-time optimal decisions. Our solutions achieve 90%+ pilot-to-production success rates and are trusted by leading companies in cement, oil & gas, power generation, chemicals, and metals industries.

Backed by industry leaders including ABB, Yokogawa, and Mitsubishi.

Learn more: www.uptimeai.com | info@uptimeai.com

Ready to see how system-level AI analysis can deliver measurable results at your facility?

Contact us to discuss your specific operational challenges and explore how AI reasoning agents can help prevent unplanned downtime, identify root causes, and scale expert knowledge across your operations.