

Beyond the Clipboard

Site Intelligence and the Future of Continuous Improvement

Better evidence for every DMAIC decision

A technical white paper for continuous improvement, Lean Six Sigma, and operational excellence leaders



Executive Summary

You cannot improve a process reliably if you cannot measure what is happening. Yet continuous improvement (CI) teams often work from short time studies, incomplete downtime codes, and records spread across multiple systems. [1,2] Those gaps make it harder to find the right problem, establish its cause, and show that a fix is still working.

At scale, those limitations affect every stage of a CI program. The consequences are practical: engineers spend time rebuilding the evidence, recurring downtime goes unexplained, and savings become harder to verify after a project closes. Better measurement should help teams make these decisions with less manual work while keeping the discipline of DMAIC (Define, Measure, Analyze, Improve, and Control).

Two evolving measurement capabilities address these gaps. Continuous computer vision creates a record of how people, vehicles, and material move through a site. Connected operational data brings production, maintenance, quality, and cost records into a shared context for AI-assisted analysis. Some problems can be answered with one; others require both. This paper uses the term Site Intelligence for the resulting measurement capability.

Site Intelligence does not replace DMAIC, standard work, measurement system analysis (MSA), statistical process control (SPC), or the judgment of a CI engineer; it strengthens the measurements behind them and supports their use at scale. Teams can compare changeovers on a consistent basis, measure how work actually happens across shifts, find root causes for recurring issues, and test the cost of different tooling decisions. Engineers still validate the measurement, establish cause, choose the countermeasure, and verify the result. The method stays rigorous; the evidence becomes easier to obtain and use.

Introduction

CI teams need to answer five practical questions: Where is time or material being wasted? How much? Why? What should change? And is the change still working?

A supervisor can investigate a single line by walking the floor and reconciling a few records. At scale, that approach becomes harder to sustain. A CI team may be responsible for multiple sites, products, shifts, and machine-paced processes; even routine observation cannot cover enough of that operation to catch every loss. Data sits across different sources, capturing different parts of the process. Engineers must spend time joining the pieces before they can investigate a problem. [2]

At Bendix Commercial Vehicle Systems LLC, an automated work cell was repeatedly striking racks and creating stoppages and cleanup work. Continuous video from Protex AI quantified more than ten labor-hours of daily cleanup, providing evidence to justify a permanent fix. The case study reports more than ten labor-hours recovered per day. [3]

This example from Bendix shows a common challenge for CI at scale: a recurring problem can cost time and money without standing out in the data used to set improvement priorities. Plant systems track cost, maintenance, throughput, quality, and production. What happens on the floor adds context: how people work, whether they follow standards, how material moves, and what happens without triggering a system record. CI teams need to bring these sources together to decide where and how to act.

This paper examines how continuous evidence supports DMAIC. The aim is to reduce the work of gathering and connecting measurements while keeping engineers responsible for interpreting the evidence, choosing the change, and owning the outcome. The following sections walk through Define, Measure, Analyze, Improve, and Control, drawing on real deployments to show what better measurement makes possible. The same evidence also supports the broader practice of kaizen. Teams need to spot problems in daily work, test changes, and check whether new standards hold; not only during a dedicated improvement event. Site Intelligence helps make those problems and results visible over time. DMAIC provides the structure for this paper, while the measurement capability supports both focused projects and everyday continuous improvement.

Where Better Evidence Changes DMAIC

Site Intelligence brings together continuous computer vision and connected operational data to strengthen the evidence behind DMAIC. It helps CI teams see what happens on the floor, connect it to plant records, and use that context at each step of the improvement process.

The two capabilities address different evidence gaps. Continuous observation helps engineers examine what happened on the floor over time. Connected operational data helps them examine the production, maintenance, and quality conditions surrounding an event. [\[2\]](#) Together, where the problem requires both, they provide a fuller basis for investigation.

The starting point varies by plant. One team may need to join siloed systems; another may already have a mature business-intelligence platform but lack physical context. A third may understand the problem and need a reliable way to monitor it. Continuous observation matters for events that would otherwise be missed. Ongoing data updates help teams spot recurring problems and emerging patterns, investigate their causes, and track whether improvements hold.

The table shows what Site Intelligence adds at each step and the decision that evidence supports.

D	M	A	I	C
Define	Measure	Analyze	Improve	Control
DMAIC STEP	MEASUREMENT CONSTRAINT	SITE INTELLIGENCE	WHAT THE TEAM CAN DO	
Define	Losses are split across MES, WMS, SAP, quality records, and spreadsheets. Building one Pareto requires manual reconciliation.	Connected operational data	Related records can be joined with little effort to surface and assess a loss across systems	
Measure	Time studies and observation cover a limited window; behavior and normal variation may not be captured across all shifts and conditions.	Continuous computer vision	Cycles and events can be captured across the study window, including periods with no observer	
Analyze	Root cause is constrained by what generated a record or happened to be visible during the gemba walk. Intermittent physical conditions can fall between those observations.	Video + correlated data	The physical condition at the time of the loss can be reviewed retrospectively and correlated to system records	
Improve	A cross-functional hypothesis can require days or weeks of manual data preparation before a decision can be tested.	Scenario modeling using connected operational data	Use interactive scenario modeling to compare countermeasures and evaluate trade-offs	
Control	Periodic checks may miss a return to previous performance.	Ongoing measurement from video and/or data systems	The same evidence can continue to measure the standard and the process after the change, making drift visible sooner.	

The measurement challenge at scale is not just a lack of data. It is a lack of context. Plant records may show downtime, defects, or added cost without explaining what was happening on the floor: which product was running, how the work was performed, whether a standard was followed, or what physical event never triggered a system record.

Connecting floor activity with plant records helps teams distinguish the loss itself from the conditions surrounding it and choose a focused investigation.

Define: Selecting and Framing the Right Problem

Define depends on knowing which problems warrant attention and setting a useful scope for the improvement work, but relying on manual downtime codes can leave major losses unexplained. Symstic reports that only 40–60% of downtime is classified when operators enter codes manually, compared with 85–95% when automated rules classify events first. [4] Pressure on the floor may help explain the gap: when a line stops, getting it running again can take priority over choosing the right downtime code. A Pareto built from incomplete or vague codes may therefore leave important downtime unexplained.

Minor stoppages also add up, but are less visible in downtime records. Nakajima included idling and minor stoppages among the major losses in Total Productive Maintenance (TPM). [5] These short, frequent events can be difficult to record manually. [1] One vendor analysis found that micro-stops account for 8–18% of planned production time across 450 instrumented sites in 30 countries. [6]

The way recorded time is grouped can also affect project scope. At one global manufacturer using Protex AI, multiple products run through shared production equipment. The facility had a 35-minute target for each changeover, but actual times ranged from 25 minutes to two and a half hours. At the time of their investigation, overrun time against that target had accumulated to 70 hours that year. The team needed to establish which differences reflected the work required and which warranted improvement. [7]

By correlating existing system records, Protex AI surfaced that the 35-minute target included changeovers with very different tooling requirements. Some required substantial tooling changes; others required none. Correlating changeover duration with the number of tooling attempts from lockout/tagout (LOTO) records showed that longer changeovers involved repeated adjustments and trials of the machine. This finding helped define the improvement effort: investigate repeated tooling attempts while accounting for the work each changeover required, and reconsider the blanket target used to judge performance. [7]

A downtime total can identify an area of concern, but defining the specific project requires examining how that time varies with the work performed. Assessing variability across shifts, products, equipment states, or operating conditions takes engineering time: records must be assembled, comparable events grouped, and possible explanations examined. Site Intelligence can make it easier to test relationships across systems, helping teams prioritize losses that only become apparent when records are examined together.

Measure: A Baseline Built From Normal Operation

A time study captures a defined period of operation. The challenge is knowing how well that period represents the process as a whole. A line may run thousands of cycles across different shifts, products, staffing levels, and operating conditions, while a study observes only a small subset.

Observation can change how people work, making the data less representative of normal operations. This concern is commonly called the Hawthorne effect, although research has not established a consistent size for it; teams therefore need to check whether their measurements reflect how work normally happens. [8]

In traditional CI methods like gemba walks, observing one non-compliant pick establishes that the standard is not always followed; measuring how often it happens, under which conditions, and how much time is associated with it requires consistent observation across the operation. Continuous measurement with Site Intelligence makes that wider coverage practical, giving the team a baseline for selecting and evaluating an improvement.

A global e-commerce and logistics company using Protex AI wanted to measure adherence to a picking procedure. The procedure required a cart when workers carried multiple packages to distant bins. Computer vision measured the work across shifts and found that 31% of picks were completed without a cart, and non-compliant picks took roughly 18 seconds longer on average. Continuous measurement with Protex AI also showed which shifts needed focused training to bring adherence back to the standard.

Site Intelligence allows Measure to produce a repeatable record of work as done across shifts and operating conditions. The same e-commerce and logistics company estimated that continuous measurement could replace approximately 13,000 hours of manual auditing each year. Reducing that observation effort would leave more capacity for the analysis and improvement work the measurements support.


Analyze: Connecting the Evidence to the Cause

Root-cause analysis depends on evidence from process records and direct observation. An intermittent physical fault may leave no useful record and may not occur while an engineer is on the floor. Data-driven methods can draw on production records to support the investigation. [9] Continuous video adds evidence of physical events that plant systems may not capture, allowing engineers to review what happened after the event.

At the Bendix work cell introduced earlier, cleanup was consuming more than ten labor-hours each day. Finding the cause through direct observation would depend on someone being present when the fault occurred. Reviewing video from periods when cleanup occurred showed the automated work cell repeatedly striking a rack. Protex AI also quantified how often the issue occurred, providing the measurement needed to quantify the loss. This gave the team a specific cause to address and a quantified loss: a defensible basis for justifying a permanent fix. [3]

Correlated data can also point to a candidate cause and a fix to test. During one investigation, the same global manufacturer introduced in Define used Protex AI to analyze downtime codes, cleaning schedules, and production data. Correlating these previously siloed datasets identified a specific product pair associated with recurring downtime when run back to back. The manufacturer introduced a cleaning step during the changeover in between these products, a change that is projected to improve yield by 1.32%. [7]

Root-cause evidence may be missing from plant records or spread across systems that capture different parts of the process. Continuous video makes physical events available for review, while correlated operational data reveals relationships between production conditions and downtime. Site Intelligence strengthens Analyze by bringing that evidence into the investigation, giving engineers a more specific explanation of the problem and a basis for choosing a countermeasure.



Improve: Testing a Candidate Change in a Working Session

Improve uses the evidence to compare countermeasures and choose a change. Here, the constraint can be the time needed to assemble a decision model and evaluate alternatives. Site Intelligence connects the relevant records and makes those comparisons practical within the time available to the team.

At one global manufacturer using Protex AI across multiple production lines, correlated data linked tooling wear over time with a recurring product defect. Replacing the tooling too early added unnecessary tooling cost and downtime; replacing it too late reduced yield. The manufacturer already had quality-inspection data showing when defects occurred. Site Intelligence connected those observations with tooling records, yield, and cost variables, and incorporated observed degradation events from the plant's inspection systems rather than relying only on the average degradation curve the customer had historically used. An auto-discovery layer from Protex AI mapped the relationships among the records, allowing engineers to model more than 40 tools and six cost variables in a working session rather than spend days simulating combinations by hand.

The model estimated that raising the yield threshold for replacing a tool from 92% to 95.5% would improve overall production yield by 2.47 percentage points and deliver about \$1.0 million in net annual value. Engineers used the model to evaluate the trade-offs and chose which specification to implement. [7]

Tool-replacement models provide an established framework for weighing replacement cost against quality loss. [10,11] The opportunity for Site Intelligence is to make those comparisons faster and better informed by current plant data. By connecting the relevant production, quality, tooling, yield, and cost data and incorporating observed degradation events, models can be updated dynamically as engineers change their assumptions. This makes a broader set of scenarios feasible to evaluate before committing to a change.

Control: Keeping Improvements on Track

Control depends on being able to verify the gain after the project closes. Continuous measurement makes it possible to monitor performance at the resolution used to establish the improvement, without repeating the observation and data-assembly work manually.

During a kaizen event, daily measurement is part of the event plan. After the event closes, monitoring is handed to the process owner under a control plan, which may use periodic checks or ongoing process monitoring. Less frequent checks can delay detection when performance slips. Research on sustaining Lean highlights why follow-up matters: achieving an improvement does not ensure it will last. [12]

For a CI portfolio, that distinction has a financial consequence. When a saving is built into the operating baseline, the budget assumes the improvement will continue. If performance deteriorates, the cost returns before the team has identified and addressed the cause. Maintaining the measurement gives the owner a more current view of whether the expected saving is still being realized.

The product-sequencing and cleaning example from the global manufacturer mentioned in the Analyze section illustrates how Site Intelligence can be used to improve the Control stage. The investigation pointed to cleaning between those products whenever that sequence occurred, rather than relying on a fixed cleaning calendar. The control question is then specific: does the cleaning take place when required, and does the associated downtime remain reduced?

The manufacturer now uses Site Intelligence to monitor the revised cleaning standard. Operational data identifies when the relevant product sequence runs, continuous video shows whether the cleaning takes place, and downtime records show whether the improvement holds. This helps the team detect when performance slips before the next scheduled audit. [7]

Site Intelligence changes how CI teams monitor and sustain gains after a project closes. The same measurement used to establish the improvement can continue to track adherence, showing whether the gain is still being realized. This also gives the CI team a more accurate basis for reporting savings. Instead of relying on the result recorded when the project closed, the team can report current performance against an established baseline.

What Remains With the CI Engineer

Site Intelligence does not remove the core responsibilities of CI. Engineers still define the problem, validate the measurement, determine whether a finding explains the problem, select the countermeasure, manage the change, and own the result. What changes is the cost, coverage, and consistency of obtaining evidence. Better measurement addresses only one part of CI; leadership, implementation, and employee involvement still matter. [13]

That distinction is important for adoption. A camera does not establish causal proof by itself. An interactive model does not make a business decision. DMAIC remains an appropriate method. Site Intelligence changes the measurement layer beneath it, making more of the operation measurable, connected, and continuously visible. Once in place, continuous measurement can support multiple improvement questions without requiring a new measurement exercise each time. Engineers still need to validate automated classifications, including false acceptances and false rejections where applicable. [14] They also need to monitor changes in the process, environment, or sensors that can affect model reliability. [15,16]

A Practical Standard for Site Intelligence

For a CI organization evaluating this approach, the useful question is whether the added measurement produces a better DMAIC output: a clearer Pareto, a more defensible baseline, a more specific root cause, a faster decision, or stronger evidence that the gain has held.

Coverage	Can the team quantify the relevant loss across shifts and operating conditions rather than only during a study window?
Context	Can an event be interpreted with the schedule, product, asset, standard, and ownership needed to make it meaningful?
Reproducibility	Are classifications and definitions consistent enough to support comparisons across lines or sites?
Auditability	Can the team show the evidence behind a finding, the decision that followed, and the measurement after the change?
Human ownership	Does the evidence support the engineer's judgment and leave the decision with the engineer?

A deployment that cannot meet those tests may still provide useful monitoring, but it has not yet demonstrated a change to the CI measurement system.

Conclusion: Instrumenting a Method That Already Works

DMAIC, Six Sigma, lean practice, MSA, gemba, and control plans remain standard operating methods for CI teams. The opportunity is to make the measurement feeding those methods less dependent on manual reconstruction and intermittent observation as the program scales.

Across Define, Measure, Analyze, Improve, and Control, the practical outputs are different: a fuller loss picture, a baseline built from normal operation, a physical cause tied to an economic consequence, a faster test of competing decisions, and evidence that the gain is still present after the project team moves on.

That is the standard against which Site Intelligence should be judged: not whether it sounds more advanced than a clipboard, but whether it gives a CI organization better evidence at the point where a better decision depends on it.

References

- [1] P. Muchiri and L. Pintelon, "Performance measurement using overall equipment effectiveness (OEE): literature review and practical application discussion," *International Journal of Production Research*, vol. 46, no. 13, pp. 3517–3535, 2008. DOI: [10.1080/00207540601142645](https://doi.org/10.1080/00207540601142645).
- [2] B. Mota, P. Faria, C. Ramos and Z. Vale, "Review of manufacturing integration between production, maintenance and quality artificial intelligence systems," *Journal of Industrial Information Integration*, vol. 47, art. 100910, 2025. DOI: [10.1016/j.jii.2025.100910](https://doi.org/10.1016/j.jii.2025.100910).
- [3] Protex AI, "Bendix Commercial Vehicle Systems Turns Data Into Everyday Safety and Operational Wins with Protex AI." [Published case study](#). Vendor-reported customer results; not an independent evaluation.
- [4] "Planned vs. Unplanned Downtime: The Data Model," Symestic. [Vendor guidance](#).
- [5] S. Nakajima, [Introduction to TPM: Total Productive Maintenance](#), Productivity Press, 1988.
- [6] "Micro-Stops on the Production Line: How to Detect and Eliminate Them," TeepTrak. [Vendor article](#).
- [7] Protex AI customer interview. October 2, 2026.
- [8] J. McCambridge, J. Witton and D. R. Elbourne, "Systematic review of the Hawthorne effect: new concepts are needed to study research participation effects," *Journal of Clinical Epidemiology*, vol. 67, no. 3, pp. 267–277, 2014. DOI: [10.1016/j.jclinepi.2013.08.015](https://doi.org/10.1016/j.jclinepi.2013.08.015).
- [9] K. Papageorgiou, T. Theodosiou, A. Rapti, E. I. Papageorgiou, N. Dimitriou, D. Tzovaras and G. Margetis, "A systematic review on machine learning methods for root cause analysis towards zero-defect manufacturing," *Frontiers in Manufacturing Technology*, vol. 2, art. 972712, 2022. DOI: [10.3389/fmtec.2022.972712](https://doi.org/10.3389/fmtec.2022.972712).
- [10] A. Jeang, "Reliable tool replacement policy for quality and cost," *European Journal of Operational Research*, vol. 108, no. 2, pp. 334–344, 1998. DOI: [10.1016/S0377-2217\(96\)00368-2](https://doi.org/10.1016/S0377-2217(96)00368-2).
- [11] P. H. Liu, V. Makis and A. K. S. Jardine, "Optimal tool replacement based on the surface roughness in finish machining," *IMA Journal of Management Mathematics*, vol. 9, no. 3, pp. 223–239, 1998. DOI: [10.1093/imaman/9.3.223](https://doi.org/10.1093/imaman/9.3.223).
- [12] B. Poksinska and D. Swartling, "From successful to sustainable Lean production – the case of a Lean Prize Award Winner," *Total Quality Management & Business Excellence*, vol. 29, no. 9–10, pp. 996–1011, 2018. DOI: [10.1080/14783363.2018.1486539](https://doi.org/10.1080/14783363.2018.1486539).
- [13] R. S. McLean, J. Antony and J. J. Dahlgaard, "Failure of Continuous Improvement initiatives in manufacturing environments: a systematic review of the evidence," *Total Quality Management & Business Excellence*, vol. 28, no. 3–4, pp. 219–237, 2017. DOI: [10.1080/14783363.2015.1063414](https://doi.org/10.1080/14783363.2015.1063414).
- [14] J. de Mast, T. P. Erdmann and W. N. van Wieringen, "Measurement System Analysis for Binary Inspection: Continuous Versus Dichotomous Measurands," *Journal of Quality Technology*, vol. 43, no. 2, pp. 99–112, 2011. DOI: [10.1080/00224065.2011.11917849](https://doi.org/10.1080/00224065.2011.11917849).
- [15] F. Hinder, V. Vaquet and B. Hammer, "One or two things we know about concept drift—a survey on monitoring in evolving environments. Part A: detecting concept drift," *Frontiers in Artificial Intelligence*, vol. 7, art. 1330257, 2024. DOI: [10.3389/frai.2024.1330257](https://doi.org/10.3389/frai.2024.1330257).
- [16] F. Hinder, V. Vaquet and B. Hammer, "One or two things we know about concept drift—a survey on monitoring in evolving environments. Part B: locating and explaining concept drift," *Frontiers in Artificial Intelligence*, vol. 7, art. 1330258, 2024. DOI: [10.3389/frai.2024.1330258](https://doi.org/10.3389/frai.2024.1330258).

*Yield improvement and annual value are model projections based on the customer's data.