PLANTWORKS · METHODOLOGY PAPER · NOVEMBER 2026

# Converting Contextualized Data into Measurable Throughput

*Zero Loss Analysis as an outcome-attribution layer for i3X deployments, and a proposed collaboration with CESMII*

Founder, Plantworks | Boulder, Colorado
trevor@plantworks.ai

> **Status of this document.** This is an unsolicited proposal prepared independently by Plantworks. It has not been reviewed, sponsored or endorsed by CESMII, and no relationship between the two organizations exists as of the date above. All characterizations of CESMII's technology are drawn from its published materials, cited in Section 11.

## 1 Executive summary

On 18 June 2026 CESMII released i3X 1.0, a vendor-agnostic open API for contextualized manufacturing information, with 44 participating technology providers, a conformance test suite and a developer toolkit. It is a serious piece of work on a problem the industry has failed to solve for thirty years.

This paper argues that i3X now faces a different problem, and that the problem is not technical. A conformance test reports whether an implementation complies with the specification. A portability test reports whether an application runs against two independently configured platforms with minimal remapping. Both are necessary. Neither produces a number that a plant's finance function recognizes as a result.

Between a conformant data layer and a measurable production gain sits a decision: which asset to act on. The i3X architecture does not supply that decision, and should not. It is a data layer, and its neutrality is a design virtue. But whoever supplies the decision determines whether the investment appears on the P&L, and the absence of that layer is, in my assessment, the largest risk to i3X adoption over the next thirty-six months.

Plantworks proposes Zero Loss Analysis as that layer. ZLA supplies three things i3X does not: a baseline referenced to physical capability rather than to an accounting standard, a method for decomposing the gap between that baseline and actual output into ranked and data-tested causes, and a conversion of each recovery into revenue or cost depending on whether the operation is production-constrained or demand-constrained. The result is an auditable throughput number with a defensible denominator.

The value runs in both directions, which is why this is a collaboration proposal rather than a vendor pitch. Plantworks' dynamic bottleneck identification requires asset identity, asset-to-process relationships, per-product rate capability and attributable loss events, continuously and across every station in a line. Under point-to-point integration that is a bespoke engineering project at every plant, which is the reason ZLA is delivered today as a consulting engagement rather than as software a plant operates for itself. A conformant i3X server is the first general route to that data shape.

Section 7 proposes four specific items of work. Section 8 proposes a four-phase pilot with stated success criteria, including the condition under which the pilot should be judged to have failed and published as such.

## 2 The attribution gap

The U.S. Department of Energy describes CESMII's expected outcomes as energy efficiency, equipment reliability, productivity gains, and improvements in safety, quality and yield. The page sets no quantified targets against any of them.[4] For an infrastructure institute at the research end of its mandate that is entirely normal. It is also a difficult position from which to address a plant manager who has been asked to justify a capital request.

The i3X 1.0 release is rigorous about the things it can measure. The Conformance Test Suite reports Full 1.0 Compliance, 1.0 Compatible, or Not Compliant.[3] Compliance is a property of software. The portability proposition, that an application built once runs on any compliant platform, is testable and worth testing. But a manufacturer who has completed a successful portability test still has an unanswered question, and it is the only question their CFO will ask: how many more cases did we ship?

THE RISK, STATED PLAINLY

CESMII's own i3X announcement warns that without consistent access to contextualized data, AI efforts stall as isolated pilots, and quotes an IDC analyst identifying brittle, point-to-point integrations on older shop-floor architectures as the main bottleneck.[3] That diagnosis is correct as far as it goes. The same fate, however, awaits a data layer that is technically excellent and operationally unattached. Adoption does not stall because the API is wrong. It stalls because nobody can name the number it moved.

This is not a criticism of the specification. It is an observation about where the specification stops, which is exactly where it was designed to stop.

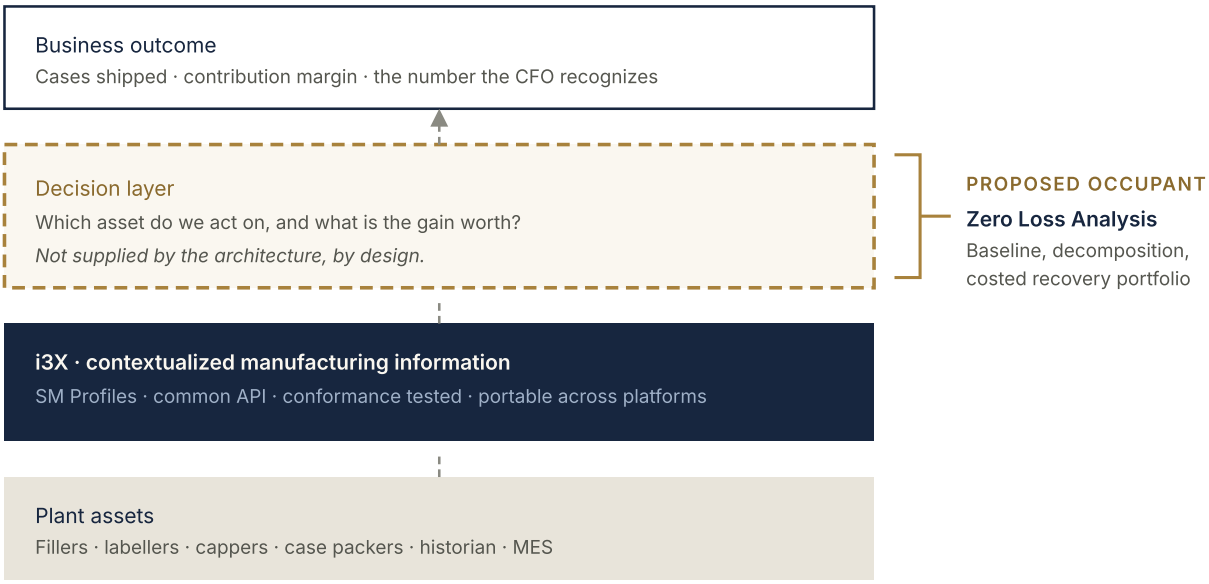



FIGURE 1 Where the specification stops. A conformant data layer carries information upward from the assets with its context intact. It does not, and is not intended to, select the asset on which action produces a plant-level gain. Whoever occupies the dashed band determines whether the investment below it reaches the outcome above it.

## 3 Why better data does not by itself produce throughput

The premise of the Theory of Constraints, originating with Eliyahu Goldratt, is that the performance of any system is governed by a single binding constraint at any given moment, and that improvement applied anywhere else produces no system-level gain.[5] The consequence for a data architecture is uncomfortable and rarely stated: contextualized data about a non-constraint is precise, well-modeled, portable information about something that cannot move the plant.

The concrete version is worth following through. A conformant data layer makes it substantially cheaper to compute overall equipment effectiveness consistently at every station in a plant. OEE is a useful measure at the constraint. As a plant-wide operating measure it is misleading, because improving it at a station that already has excess capacity pro-

duces no additional shipments. A superior data layer therefore makes it cheaper, faster and more consistent to industrialize a measurement error that is already widespread.

I have watched this pattern for forty-seven years, under several different names. Kaizen events on a station that already had idle time. Six Sigma on a defect that was never the limit. A warehouse reorganization that made the plant look dramatically better and shipped exactly the same volume. Different methods, same mistake. The methods were not defective. The targeting was. My concern is that a data layer good enough to make targeting cheap will also make mistargeting cheap, and at greater scale.

The honest conclusion is narrow: the data layer is not self-directing, and it was never supposed to be. The question is who supplies the direction.

---

## 4 What Zero Loss Analysis supplies

Zero Loss Analysis is a quantification framework I introduced in the *Zero to Hero* working paper while a director at Deloitte in 2009, in the firm's Strategy & Operations practice. The paper was subsequently republished in the Operational Excellence Society Body of Knowledge and remains in circulation.[6] Four of its elements matter to the argument here.

### A baseline referenced to physics, not to an accounting standard

*Full Potential* is the output an operation would produce if the binding constraint ran at its maximum demonstrated sustained capability for the entire production window, with no downtime, no slow running and no waste. *Full Process Efficiency* is Optimum Process Hours divided by Actual Process Hours: the proportion of run time that produced first-quality output rather than one of the three loss categories.

The distinction that matters is the denominator. A standard rate is negotiated, and once negotiated it absorbs the losses that were present when it was set. A maximum demonstrated sustained capability is observed. An efficiency figure referenced to it cannot be improved without real recovery, which makes it suitable as the baseline in an outcome claim that a third party may audit.

### A decomposition that is tested rather than brainstormed

Losses at the constraint fall into three categories, in this order: **downtime**, **slow running**, and **waste**. Slow running is production lost because the constraint runs below capability while still running. It appears in neither the downtime log nor the quality log, and in my experience it is usually the largest of the three once analysis is complete.

*Variance tree analysis* decomposes the total loss at the constraint into component causes, ranked by magnitude and tested against historian data rather than generated in a workshop. This is the point of contact with a data layer: the method is only as good as the data it can interrogate, and it is structurally hostile to causes that cannot be tested.

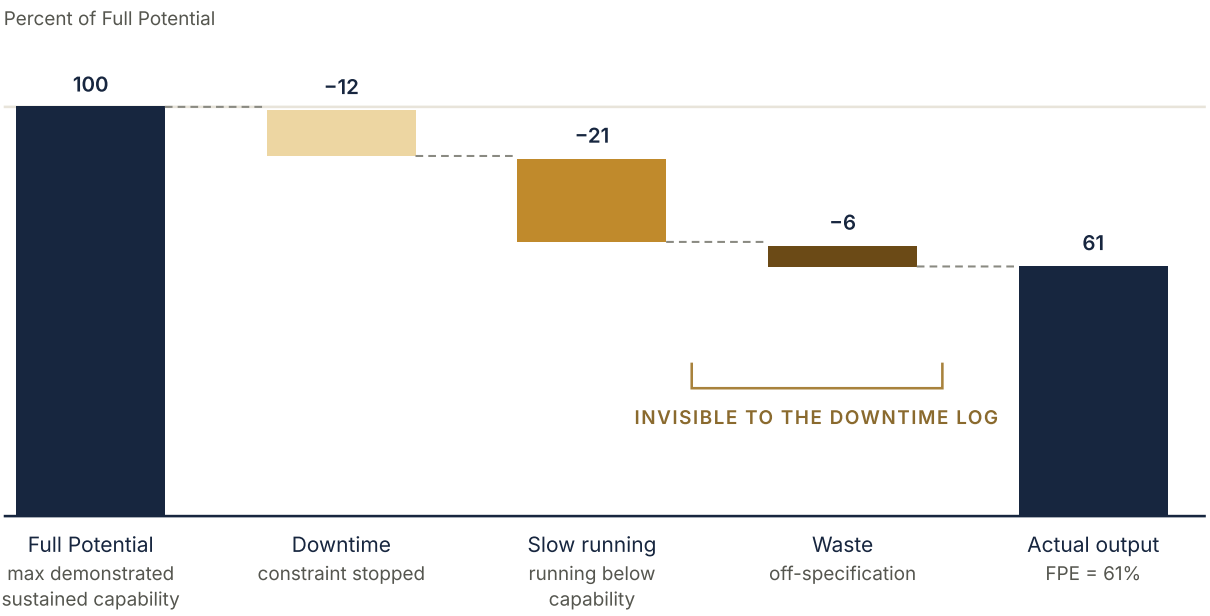



FIGURE 2 The shape of a variance tree at the binding constraint. Loss is measured against observed physical capability rather than a negotiated standard rate, which is what makes the residual auditable. Slow running is typically the largest of the three categories and appears in neither the downtime log nor the quality log. *Proportions are illustrative of the usual pattern, not measurements from a specific plant.*

### A conversion that finance recognizes

Recovered capacity is worth different things in different market conditions. In a *production-constrained* operation, where the market would buy more than the plant can supply, recovery converts directly to incremental revenue at contribution margin. In a *demand-constrained* operation it converts to cost reduction. Stating which condition applies is what separates a throughput claim from a savings estimate, and most improvement programs never make the distinction.

> WHAT THIS OFFERS CESMII SPECIFICALLY
>
> An attribution method. ZLA answers "how much, and compared to what," with a baseline derived from observed physical capability and a decomposition tested against the plant's own historian. That is the shape of claim an institute accountable to a federal funding agency can defend, and it is not a claim the data layer can generate on its own.

## 5 Why dynamic bottleneck identification needs i3X

The reciprocal argument is stronger than the first one, and it is the reason I am writing to CESMII rather than about it.

Conventional constraint practice treats the bottleneck as a location: a one-off study identifies a station, and the plant works on it. In discrete operations running multiple products, that is wrong often enough to be dangerous. Each product runs at a different rate through each station, a pattern I call *speed curves per SKU*, so the binding station for a given week is a function of that week's product mix. A plant can also hold a *hidden queue*: a station at moderate utilization, typically sixty to eighty percent, that is nevertheless binding throughput because its effective capacity is degraded by losses the measurement system does not surface. The station everyone complains about is frequently not the constraint.

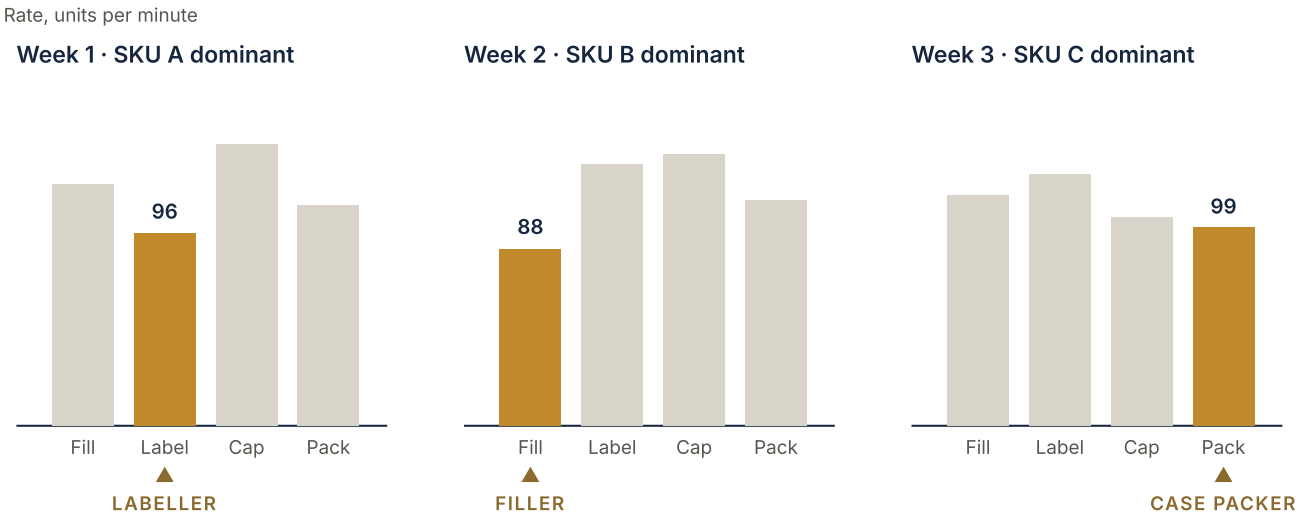



FIGURE 3 Same line, same equipment, three weeks. Because each product runs at a different rate through each station, the binding station follows the product mix: the labeller in week one, the filler in week two, the case packer in week three. Brass marks the binding station; gray marks stations with capacity to spare. A one-off constraint study would have found one of these three and been wrong for two-thirds of the quarter. *Rates are illustrative of the pattern, not measurements from a specific plant.*

*Dynamic bottleneck identification* is therefore a computation, not a study: which station is currently binding throughput, given the mix being run and the actual state of the operation, recomputed each shift. Doing that continuously requires, across every asset in the line:

- asset identity, stable across systems;
- asset-to-process relationships, so a station's position in the flow is known rather than inferred;
- product identity on the running order;
- rate capability per product per asset;
- loss events attributable to a specific asset and classifiable into downtime, slow running or waste.

Assembled from point-to-point integrations, that is a bespoke engineering project at every plant, re-done whenever a line changes. It is the brittleness CESMII's own release identifies.[3] It is also the specific reason ZLA reaches plants today as a consulting engagement rather than as software a plant runs for itself: the method is portable, and its data plumbing is not.

A conformant i3X server is the first general route to that data shape I am aware of. The effect on ZLA is not a cost saving but a change in delivery model.

## A candidate profile contribution

If the binding station depends on product mix, then rate capability per product is not ancillary information. It is a first-class attribute of an asset, and it belongs in the information model rather than in a spreadsheet held by a process engineer. I do not know whether CESMII's existing SM Profiles already model per-product rate capability. If they do, I would like to use it. If they do not, Plantworks would contribute the model.

## 6 An anchor application, not another consumer

Standards are adopted when an application people want cannot be built without them. Reduced integration cost is a real benefit and a weak adoption argument, because it is a cost-side claim competing against the status quo of integrations that already exist and already work.

Dynamic bottleneck identification across a multi-asset line with changing mix is in the other category. It is not merely cheaper with a common contextualized API; at scale, and re-computed every shift, it is impractical without one. An application of that kind gives i3X an adoption argument denominated in cases shipped rather than in engineer-hours avoided.

I note also that the 1.0 release includes an open-source MCP server for natural-language queries against a connected i3X server.[3] The direction of travel there and the direction of my own current work are the same. A language model reasoning about a plant's constraint needs to know what a tag means, which asset it belongs to and how that asset sits in the process. Without that, every deployment is a hand-built mapping exercise, which is the ceiling any constraint-reasoning application will hit as it scales.

## 7 Proposed collaboration

Four items, in descending order of how much I think each is worth to CESMII.

PROPOSED ITEMS OF WORK

| Item | What it involves | Value to CESMII |
|---|---|---|
| **1. Reference application** | Plantworks builds its variance tree analysis tooling as an i3X client and runs it at a pilot site against two independently configured compliant platforms. | Demonstrates the portability proposition with an application whose output is a production number, not a remapping count. |
| **2. Outcome attribution** | A jointly reviewed method for stating deployment results in throughput and margin terms, with Full Potential and FPE as the baseline convention. | Supplies a defensible answer to "what did it produce," usable across the ecosystem and in reporting to DOE. |
| **3. Profile contribution** | Per-product rate capability as an asset attribute in the SM Profile library, subject to the open question in Section 5. | Extends the profile library into mix-dependent capacity, which mix-driven discrete operations need and which I have not found modeled elsewhere. |
| **4. Curriculum adjacency** | Mutual cross-referencing between CESMII's certification programs and Plantworks Academy. | CESMII's curriculum covers architecture and information modeling; mine covers constraint identification and loss decomposition. The content does not overlap and the audience is identical. |

## 8 Proposed pilot design

One discrete manufacturing line, food and beverage or CPG, running at least four products with changeovers occurring within a shift. Mix variability is the condition that makes the exercise meaningful; a single-product line will not test the proposition.

FOUR PHASES AND THEIR SUCCESS CRITERIA

| Phase | Duration | Deliverable | Pass condition |
|---|---|---|---|
| **1. Baseline** | ~2 weeks | Full Potential and FPE for the line, computed from historian data accessed through i3X, with all assumptions stated. | Baseline reproducible by a third party from the same data and stated assumptions. |
| **2. Decomposition** | ~3 weeks | Variance tree: constraint loss decomposed into downtime, slow running and waste, with causes ranked by magnitude. | Each ranked cause tested against plant data, not asserted. Causes that cannot be tested are reported as untested. |
| **3. Portability** | ~2 weeks | The same application, unmodified, run against a second independently configured compliant platform. | Remapping effort measured in engineer-hours and published, whatever the figure. |
| **4. Exploitation** | 90 days | Top-ranked recoveries acted on without capital expenditure; change in output measured against the Phase 1 baseline. | A stated change in cases shipped, attributed against the baseline, with the production- or demand-constrained condition declared. |

> FAILURE CONDITION
>
> If Phase 4 shows no measurable throughput change against a correctly established Phase 1 baseline, that is a real result about the method and should be published as one. I would rather enter this with the failure condition agreed in advance than argue about attribution afterwards. A pilot that cannot fail proves nothing, and an institute accountable for public funds should not be asked to accept one.

## 9 About Plantworks

Plantworks is a manufacturing methodology practice based in Boulder, Colorado, focused on discrete operations, principally food and beverage and consumer packaged goods. It is built around Zero Loss Analysis and the Theory of Constraints.

I have spent forty-seven years in manufacturing operations across automotive, consumer packaged goods, food and beverage, pulp and paper, chemicals, and oil and gas. I have held senior roles at Schneider Electric, Invensys and Deloitte, and served as chief executive of Hagen & Company. I wrote the *Zero to Hero* working paper at Deloitte in 2009, introducing Zero Loss Analysis to the firm's global Strategy & Operations practice; it was later republished in the Operational Excellence Society Body of Knowledge. My book *Where's Your Herbie?*, a Theory of Constraints treatment written for people who run discrete plants rather than for a seminar room, was published in 2026.[7]

I also work inside open standards for industrial automation. I hold a commercial O-PAS license from The Open Group, co-chair Marketing and Outreach for the Open Process Automation Forum, and served as a Forum co-chair from 2017 to

2023. I mention it because it shapes how I have written this paper: I understand what it takes to get a specification agreed, conformance tested and adopted, and I am not under the impression that the hard part of i3X is behind you.

## 10 Limitations and open questions

Stated directly, because a proposal that hides its weak points wastes the reader's time.

- **The software is in development.** The variance tree analysis tooling described here is in active build and is not deployed at scale. The pilot in Section 8 would be an early deployment, and should be scoped as one.
- **Recovery ranges are practitioner observations.** Figures I have cited elsewhere for recoverable throughput in the first ninety days come from my own engagements over several decades. They are not a published dataset, a controlled study, or a statistically defensible sample, and I do not present them as one. The pilot is proposed partly because I would like a measured number rather than an observed one.
- **I do not know whether per-product rate capability is already modeled** in CESMII's SM Profile library. Section 5's offer is conditional on the answer.
- **I could not confirm the relationship between i3X and OPC UA or MQTT.** i3X is commonly described as complementing both. The 18 June 2026 release materials I reviewed do not address either protocol,[2][3] so I have made no claim about it here and would welcome correction.
- **Prior art within the ecosystem is unknown to me.** If a participant among the 44 technology providers is already pairing contextualized data with constraint-based prioritization, that materially changes this proposal, and I would rather be told than discover it later.
- **There is no existing relationship.** This document is unsolicited and carries no endorsement. Nothing in it should be read as implying that CESMII has evaluated Zero Loss Analysis or Plantworks.

## 11 References


1. U.S. Department of Energy, Advanced Materials & Manufacturing Technologies Office, "Collaborative Ecosystems Smart Manufacturing Innovation Institute (CESMII)." energy.gov/eere/ammto/collaborative-ecosystems-smart-manufacturing-innovation-institute-cesmii (accessed 4 November 2026).
2. Manufacturing USA, "CESMII releases i3X 1.0: complete, implementation-ready open API specification for manufacturing." manufacturingusa.com (accessed 4 November 2026).
3. CESMII, "CESMII Releases i3X 1.0," 18 June 2026. cesmii.org/cesmii-releases-i3x-1-0 (accessed 4 November 2026). Source for the release date, the 44 participating technology providers, the contents of the release package including the Conformance Test Suite and MCP server, the "API chaos" characterization, and the quoted IDC assessment.
4. DOE, as above. Source for the institute's expected outcomes and the absence of quantified targets against them.
5. Goldratt, E. M. *The Goal.* North River Press, 1984. Origin of the Theory of Constraints and of the hiker, Herbie, used as the name for the binding constraint.
6. Cusworth, T. *Zero to Hero.* Working paper, Deloitte, 2009. Republished in the Operational Excellence Society Body of Knowledge.
7. Cusworth, T. *Where's Your Herbie?* First edition, independently published, Boulder, Colorado, 2026.

Plantworks · Boulder, Colorado · trevor@plantworks.ai · plantworks.ai
Prepared 4 November 2026. Independent proposal; not endorsed by or affiliated with CESMII, the U.S. Department of Energy, or Manufacturing USA.