ACTEMIUM

Empowering Manufacturers to Unlock the Full Value of Their Data

INTRODUCING

ImpactNOW™

TURN UNUSED DATA TO ROI IN WEEKS

A rapid, low-risk entry to AI that helps manufacturers uncover hidden gains – without waiting for a full digital transformation.

WHAT ImpactNOW™ DELIVERS

SPEED

Value in 3-5 weeks, not months

CLARITY

Clear ROI case for optimization

CONFIDENCE

A plan backed by a trusted team that works in your best interest.

How It Works

- On-Site Discovery
- Data Extraction
- Exploratory ML Modeling
- ROI Case Report
- Optimization Path Forward

Key Deliverables

- Comprehensive **ROI Case Report** with Quantified savings and KPIs.
- Initial **ML/AI insights** highlighting critical patterns, correlations, and causation.
- A strategic **path forward** with clear expectations for success

Ready to uncover ROI in your data? Let's start your ImpactNOW™ discovery today.

ACTEMIUM

CASE STUDY

Unlocking Hidden Factory Value with ImpactNOW and Automated Startup Control

THE PROBLEM

Across modern manufacturing, as much as 80% of operational technology (OT) data goes unused. Plants are equipped with thousands of sensors and control loops producing valuable process information every second, yet most of it remains disconnected from meaningful analysis or decision-making. This lack of visibility limits both operators and leadership — decisions are made based on experience rather than insight, and inefficiencies accumulate quietly, invisible to financial reporting.

For our Customer, this challenge became clear when we were initially asked to address PID loop instability during production upsets. However, when we began exploring the data through Actemium's ImpactNOW methodology, it became evident that the core issue wasn't the loop tuning itself — it was what was happening before the loops even stabilized. ImpactNOW's structured EDA workflow revealed that the yogurt separator startup process was highly manual and inconsistent. Operators entered flow setpoints, monitored conditions, and decided when to switch the PID from flow-based to protein-based control purely by judgment and tribal knowledge.

While the perception was that startups took about seven minutes, data showed they averaged 13–14 minutes, with each run producing 390–420 gallons of off-spec product. In nearly one out of every five runs, systems were accidentally left in manual flow mode for half the batch, leading to in some cases product holds, inconsistent quality, and unmeasured financial waste — a classic example of the "Hidden Factory" that ImpactNOW is designed to expose.

THE SOLUTION

Using a combination of Exploratory Data Analysis (EDA) and Machine Learning (ML), following the structured phases of Actemium's ImpactNOW approach, our team approached the problem differently. We didn't start by just tuning the PID loops as directed — we started by letting the data reveal the true source of variation.

By analyzing hundreds of historical startup runs, ML models revealed consistent patterns: some operators achieved stable, in-spec conditions in as little as six minutes. These repeatable "golden signatures" showed that faster, more efficient startups weren't random — they were predictable and therefore could be automated.

With those insights, we developed an Auto Sequencing algorithm that detects when product quality has stabilized and automatically transitions the separator from flow to protein control. This eliminated manual judgment, standardized the startup process, and provided operators and engineers with clear visibility into every step of the transition. The solution went live on the first separator in June 2025, and performance data immediately validated the accuracy and repeatability with measurable ROI.

THE BENEFITS

The results were measurable and immediate. The automated startup sequence reduced startup waste by 3.7 minutes per run — roughly 132 gallons of product saved each cycle. By reducing the time between startup and steady-state operation, the system not only minimized material loss but also stabilized the downstream PID loops, improving consistency and reducing operator workload.

Perhaps even more importantly, the Customer gained data-driven visibility into a part of the process that had never been measured before — exactly the type of impact ImpactNOW targets. The same analytics and ML framework that uncovered this opportunity can now be reused across other process areas, giving a replicable toolkit for finding and quantifying hidden factory losses across its network.

THE ROI

The financial impact of the Auto Sequencing implementation is substantial. Across 16 separators, the verified reduction in waste equates to a total **annual savings of approximately $2 million in material costs** alone. Each startup run now avoids roughly $264 in material waste per separator, and the material cost of waste per minute has dropped significantly.

The payback period for this initiative was less than three months, and the improvement is sustainable because it's built directly into the control system — not dependent on operator behavior.

THE NEXT STEPS

The success of the Auto Sequencing rollout marks the beginning, not the end, of this Customers digital transformation journey. The next phase focuses on predictive process control, where the same ML techniques will be used to forecast protein yield based on upstream variability and automatically adjust setpoints before quality is impacted.

By integrating these predictive algorithms directly into the PLC control layer, the process will evolve from reactive control to proactive optimization

In short, the project that began as a PID tuning request became a $2 million/year yield recovery initiative, powered by EDA, Machine Learning, Automation, and the structured ImpactNOW methodology — proving the value of turning unused data into operational and financial excellence.

ACTEMIUM

CASE STUDY

ImpactNOW Data-Led Optimization of High-Speed Filling Operations

THE PROBLEM

Modern filling systems generate large volumes of operational data, yet very little of it is used to quantify or control head-level variation. As a result, mechanical drift, poor-performing heads, and hidden giveaway accumulate silently over time. For this Customer, inconsistent fill weights were believed to be an unavoidable part of the process. However, once Actemium applied its ImpactNOW methodology — a structured workflow for rapidly transforming OT data into operational insight — the core issue became clear: the filler's 32 heads were not behaving uniformly, and the system had no built-in means to detect or compensate for the divergence.

Through data normalization, per-head variance analysis, and model-driven diagnostics, ImpactNOW revealed that several heads exhibited elevated variability well beyond what the system could achieve, with one head in particular showing instability severe enough to suggest imminent mechanical failure. These losses were occurring continually but were not visible to operators or engineering teams — a textbook example of hidden process waste that ImpactNOW is designed to expose.

THE SOLUTION

Using Actemium's ImpactNOW methodology, the team began by transforming a raw, unstructured long-format historian export into an engineered dataset capable of supporting head-level diagnostics. This required reconstructing the filler's operating cycles, normalizing timestamps, and pivoting more than one hundred tags into per-head feature vectors. Once aligned, the data exposed the true mechanical behavior of the system: variations in filling time, inflight weight, product density, inlet flow, temperature, and pressures — all key drivers of final filling mass. The engineered dataset provided the first high-resolution, quantitative view of head-to-head performance, revealing patterns, drift, and instability that had previously been invisible to operators.

With clean and structured inputs, the team trained individual machine learning models for each of the 32 filler heads. These models used nine engineered features to predict final filling weight with high fidelity, achieving R^2 values consistently above 0.90 for most heads. Beyond raw accuracy, model explainability quantified how each head responded to process conditions, enabling clear differentiation between normal mechanical variation and abnormal or failing behavior. The models also established the statistically achievable "gold standard" variation based on the best-performing heads — creating a benchmark for what the system could accomplish uniformly. This combination of engineered data, per-head prediction models, and interpretable variability baselines formed the foundation for maintenance prioritization, predictive alerts, and the feasibility assessment for in-run model-based recommendations.

THE BENEFITS

The ImpactNOW approach produced measurable insight in days rather than months — with immediate financial impact.

Immediate Impact:

- Identification of a filler head as a malfunctioning component
- Immediate maintenance action prevents future downtime and reduces giveaway.
- Clear line-of-sight to achievable CV reduction
- If all heads align with best-performers, variation can be reduced by 15–30%, directly reducing underfills and giveaway.

Strategic Benefit:

Perhaps most importantly, the Customer now has:

- A validated, data-driven understanding of head-level mechanical behavior
- A foundation for predictive maintenance
- A blueprint for closed-loop optimization
- A scalable methodology that can be repeated across every filler in the network

This is the core promise of ImpactNOW: turning unused OT data into measurable operational and financial results.

THE ROI

A focused overfill/underfill analysis across 5.1M fills shows the filler is currently generating ≈$1.0–$1.1M in annual net overfill cost. Because the ML models demonstrate that many heads can consistently achieve far tighter distributions, a **4–8-gram reduction in average overfill** is realistically achievable through per-head variation reduction and drift detection. This translates to **$300k–$800k in yearly savings** on this filler alone.

Additionally, correcting the mechanically drifting head provides an immediate $40k–$50k benefit while reducing underfill risk.

THE NEXT STEPS

The next phase focuses on moving the analysis into real-time operations. We will deploy an edge layer that streams high-frequency process data directly from the filler, runs the per-head machine learning models on the line, and generates immediate trim adjustments to reduce variation. This same layer will monitor each head for behavioral drift and trigger early maintenance alerts when performance begins to degrade.

The architecture is designed to scale, allowing the same workflow—data acquisition, per-head modeling, and real-time optimization—to be replicated across additional fillers and business units, creating a consistent and repeatable method for reducing overfill and improving process capability enterprise-wide.