

MULTI-VARIATE CYCLE INTELLIGENCE — WHAT YOUR MES AND HISTORIAN CANNOT TELL YOU

Save the cycle. Improve yield.

Every plasma fractionator runs cycles that drift weeks before final QC sees it, investigates the same root cause across campaigns and plants without connecting it, writes CAPAs for events that were preventable, and never quite reconciles the dollar impact of yield loss to a number a CFO can defend. **These look like four problems. They are four readings of one signal** — the trajectory of the fractionation cycle itself. And unlike every other modality, the raw material cannot be re-ordered.

01 Drift starts weeks before final QC sees it

02 One root cause, hidden across cycles and plants

03 CAPAs written for what could have been prevented

04 Yield costs absorbed, uncouned

The data is already in your plants

You don't need new sensors, new instrumentation, or a re-platforming project. The signal that predicts each cycle's outcome exists days to weeks before the cycle tells you — sitting in the same systems you already run, across every fractionation site.

PI HISTORIAN · EVERY PLANT

Density, complete

Every fractionation cycle generates thousands of parameter reads per hour — Cohn temperature control, Fr.II+III pH, ethanol dose and addition rate, agitation, hold-step conditions, chromatography column cycles. Structured, high-frequency, already logged.

MES / EBR (PAS-X, WERUM) · EVERY CYCLE

Genealogy captured

Pool composition, donation genealogy, source-versus-recovered ratio, pool age, raw-material lots, operators, equipment IDs, deviations, exceptions. Every cycle's full lineage structured and permanent in your batch-record system.

LIMS · EVERY RELEASE

Outcomes measured

Aggregate content, HCP levels, potency against release spec, IgG yield, purity, subclass distribution, stability. The dependent variable of every cycle — waiting to be joined to what upstream drove it.

SENIOR OPERATORS · EVERY PLANT

Signal held in heads

Which ethanol supplier lots to distrust. Which centrifuges run warm on second shift. Which pool profiles precede an aggregate creep. The pattern lives in the memory of your most experienced fractionation operators — carried in conversation, not captured in any system, and lost every time one of them retires.

Read it against your own plants

If you run **PI Historian, MES / eBR, and LIMS** — and you do — you already have three of the four inputs the platform reads. The fourth lives in the heads of your senior operators. What you don't have is a layer that reads them together, in real time, against the pattern of your own history, across your plant network. The signal is not missing. It is scattered — across three systems and the memory of the people who built your plants.

02 — WHY PLASMA IS NOT LIKE THE REST

You cannot **re-order the raw material**

Every modality loses batches. Plasma is the only one where the input is a donated, irreplaceable human material — and where a single pool feeds four to six commercial products at once. That changes what a lost cycle actually costs.

A small-molecule campaign can be re-run next week from the same API. A biologics campaign can be re-seeded from the same working cell bank. A fractionation cycle cannot be re-run, because the pool that fed it does not exist twice. The donations behind it were collected across weeks and centers, screened, tested, held, and committed. When a cycle is lost, what is written off is not inventory — it is donor material that cannot be replaced at any price, plus the held capacity, plus the downstream schedule across the whole cascade.

One pool. One cascade. Four to six products. A single upstream drift does not cost you a batch — it costs you a portfolio.

That is the structural asymmetry. Because the Cohn cascade is sequential, an intake or Fr.I decision propagates into IVIG, SCIG, albumin, and the hyperimmunes simultaneously. The organizations that fractionate plasma are, by definition, running the modality with the least tolerance for late detection — and are largely monitoring it with the same single-parameter tooling as everyone else.

03 — THE COST IS REAL, AND MOSTLY INVISIBLE

One drift, four department problems

The loss is never labeled “loss.” It shows up as a routine investigation, a lot on hold, a potency result at the bottom of spec, a delayed release. Here is the same underlying cycle drift, surfacing as four separate departments’ problems.

MSAT

The recurring investigation

A subtle upstream drift — an ethanol lot at the edge of the historical envelope, a pool-age shift, a centrifuge running warm — produces a downstream deviation. Root-cause takes six to twelve weeks. The same investigation repeats next campaign, same conclusion, same delayed action.

Owned by: manufacturing sciences

The CAPA that arrives too late

CAPA compilation begins after the deviation is confirmed and the cycle is already at risk or lost. Documentation is thorough. The cycle it would have saved is gone — along with the pool behind it.

Owned by: QA / QC

The cycle written off

Aggregate breakthrough, HCP carry-over, or potency sliding below the lower release spec — caught at final QC. Fully-loaded cycle cost absorbed, held capacity absorbed, downstream schedule absorbed across the cascade. The signal was in the trajectory one to three weeks earlier.

Owned by: operations

The supplier drift that reached your cascade

An ethanol (or resin, or filter, or buffer salt) supplier's CoA passes incoming QC — technically within spec, but sitting at the edge of the historical envelope. No automated alert fires. The lot moves into cold-ethanol fractionation at two sites. Several cycles later, the pattern connects — and by then the pools are committed. The signal was in the supplier's own release data and your incoming-QC history all along. Nobody read them together against your cycle outcomes until the deviation surfaced.

Owned by: procurement & incoming QC

The cycle is one process. You see many.

This is the core of it. A fractionation cycle is a single, continuous multi-parameter process running for weeks from intake to release. Your organization is built to experience that process as a set of disconnected signals, each captured by a different system and routed to a different team.

HOW THE CYCLE EXPERIENCES IT

One continuous multi-parameter trajectory.

Every parameter connected to every other. Every hour connected to the outcome three weeks later. To the cycle it is one process, one story, drifting or holding across dozens of dimensions simultaneously — pool age, source-versus-recovered ratio, Fr.II+III pH, ethanol dose and lot signature, Cohn temperature control, chromatography column cycles, viral-inactivation hold, incoming raw-material lot, plant, and operator.

HOW YOUR SYSTEMS EXPERIENCE IT

Four disconnected signals.

The same trajectory is split at the door and handed to teams who never compare notes:

MES / PAS-X sees a batch record.

PI Historian sees a trend.

LIMS sees a QC result.

MSAT sees a deviation.

Every cycle has multiple stakeholders across many systems — and the industry has quietly accepted the yield drift that follows as the cost of doing business.

Every system in that list is doing its job. MES records the cycle. Historian trends the parameters. LIMS releases the outcome. Manufacturing sciences investigates the deviation. None of them is broken — and none of them is looking at the cycle as one continuous trajectory across the cascade and across sites. The next section is what that costs.

05 — THE FOUR SYMPTOMS

What the split costs you

Each disconnected signal becomes its own chronic symptom — and each maps to value that was recoverable, if anyone had read the trajectory in time.

01

Cycles that fail at release

The cycle drifted weeks before final QC flagged it — pool composition, Fr.II+III conditions, ethanol lot. The signal was in the trajectory. It was just unread until the release panel came back, and by then the cycle was already committed.

02

The same investigation, three campaigns in a row

The root cause is found at campaign one, forgotten by campaign two, rediscovered at campaign three. Each investigation runs six to twelve weeks before the next cycle can start. Each was preventable at the second occurrence.

03

CAPAs written for what could have been prevented

04

Yield variability absorbed as the cost of doing business

The CAPA is thorough. It documents what happened, why, and what will change. It arrives after the cycle is lost. The best CAPA is the one you never had to write.

Campaign-to-campaign IgG-yield swings treated as inherent to fractionation. The pattern of what drives the high-yield cycles is sitting in your own historical data — read once and never revisited.

06 — WHY THE CURRENT MODEL KEEPS LOSING GROUND

The gap widens on its own

None of this is a failure of effort. Process scientists work hard inside the tools they have. The model itself is what keeps losing ground — for four compounding reasons.

01 You find out too late

Conventional MES and Historian reporting is a rear-view mirror. The deviation is flagged after it happens. The release panel arrives after the cycle is committed — and in fractionation, “committed” means the pool is already consumed. By the time the number lands on a dashboard, the moment to act on it has passed.

02 You pay four times for one signal

Four teams, four tools, four budgets — each chasing a different reading of the same underlying cycle trajectory, none of them aware of the others. Across every plant and every product in the cascade. The spend is real and recurring; the coordination is not.

03 Single-parameter monitoring cannot explain the variance

Specifications are set and monitored one parameter at a time. Most of the cycle-to-cycle variability that matters is interactive — an ethanol lot signature that is harmless at one pool age and consequential at another. One parameter at a time cannot see an interaction, and no operator can hold sixty dimensions across three plants at once.

04 You forfeit upside within reach

Higher IgG yields left on the table. Drifting cycles lost that could have been preserved. Faster campaigns delayed because a signal was read at week three instead of week one. The upside was there — the mechanism to capture it wasn't.

07 — ONE SIGNAL, FOUR RECOVERIES

Read the trajectory **once**

Here is the shift. The same fractionation trajectory — read once, from data you already own — surfaces across all four problems at the same time. Not four tools. One reading, four uses, entirely within your own systems.

THE UNIFIED CYCLE SIGNAL

**One cycle, surfaced four ways
the moment its trajectory
shifts**

RECOVERY 01 • PREDICTION

**Cross-signal early
warning**

RECOVERY 02 • BLUEPRINT

**Gold Standard, reverse-
engineered**

Ethanol-lot signature, pool age, Fr.II+III pH, Cohn temperature, chromatography cycles, and downstream QC read simultaneously. The pattern flags drift one to three weeks before release testing — before the deviation is visible to any single system.

Your best cycles, reverse-engineered — specific parameter targets, timing windows, thresholds, ranked by effect size. Not “watch Fr.II+III pH.” Rather: “hold pH in range X during precipitation; more than two excursions per cycle predicts aggregate above Y.”

RECOVERY 03 · PREVENTION

Prevention over documentation

Ranked corrective actions delivered during the run — which parameter to correct, by how much, in what window — not CAPAs assembled after the cycle is lost. The best CAPA is the one you never had to write.

RECOVERY 04 · NETWORK

Root cause, across your network

One drift signal, connected across cycles and sites. One ethanol lot, two products, two plants — surfaced as one pattern, not several separate investigations. The connection a single-plant, single-system view genuinely cannot see.

Today you pay four teams to chase four versions of this signal, and none of them sees the other three. **The recovery is not new instrumentation. It is reading the trajectory you already have — once, and together, across every plant and every product in the cascade.**

One more thing. The same trajectory, read one degree differently, starts to question things the industry has treated as fixed — how pools are composed, how potency distributes, which upstream signals are worth reaching for outside your own four walls. That’s not a claim we make in a white paper. It’s a conversation we’d love to have.

A wrapper around your systems, not a replacement

The first question a validated environment asks is what this does to the validated environment. The answer is nothing. FyndEm Quality reads from your existing systems and never writes back.

01

Reads existing data

Connects to PI Historian, MES / eBR (PAS-X, Werum), LIMS, and DCS by file export or direct connector. No new sensors, no new hardware, no changes to current data flows.

02

Never writes back

Zero mutations to any source system. Validated batch records, MES workflows, and process control loops continue exactly as they do today. Read-and-advise only.

03

Decision support, not decision making

Intelligence surfaces through a separate interface. Your process scientists and quality teams act on it or don't, at their discretion, inside your existing procedures.

04

Deterministic, reproducible, explainable

The production decision path is pure mathematics and statistics — no AI or LLM in the live path. The same inputs produce the same outputs, and every alert can be traced to the parameters that generated it. Aligned to the ICH Q9 / Q10, 21 CFR 640, and Ph. Eur. 3.5 framework your quality system is already built on.

09 — WHAT IT'S WORTH TO YOU

A number from data you already run

The natural question is “what is this worth for us, specifically?” — and you can get a defensible answer before any engagement, any contract, or any raw data leaving your building.

Your estimate comes from your own history

Using 18–24 months of your Historian, MES / eBR, and LIMS data — under a data-use agreement, with de-identification handled up-front — the four recoveries can be sized into a conservative Year-1 range, decomposed by recovery stream and by plant, with every assumption stated. No new sensors. No MES disruption. Just your own cycle history, read the way the four-recovery model reads it.

STEP 01

Start from your cycles

18–24 months of Historian + MES + LIMS supplies pool genealogy, upstream signatures, release outcomes, and deviations — the full record across your fractionation network.

STEP 02

Apply conservative rates

Each recovery stream is sized with deliberately conservative, fully disclosed realization assumptions — your own fully-loaded cycle cost, your own investigation cost, your own held-capacity impact.

STEP 03

See your range

You receive a Year-1 recoverable range for your specific plant footprint — decomposed by recovery stream and by product in the cascade, every assumption disclosed.

All quantifications are produced individually — from your own historical data and your own cost drivers, under a data-use agreement. Nothing in this paper is presented as your number.

THE NEXT STEP

Reach out to us today

YOUR CENTROIDAI CONTACT

Raj Subramanyam

Founder & CEO, CentroidAI · raj@centroid-ai.com

Contact Us

FyndEm Quality — Plasma is one variant within the broader **FyndEm Quality** suite. **Solid Dose** covers small-molecule and biosimilar manufacturing, including API; **Biologics** covers mammalian and microbial bioprocessing; **Corrugated** covers containerboard and box-plant converting. One platform, one architecture, tuned per industry and modality.

CentroidAI · **FyndEm Quality — Plasma**

Industry references: BioPlan / BioProcess annual manufacturing surveys (2008–2024, biologics — indicative, not plasma-specific); PDA and ISPE process-performance benchmarks; ICH Q9 / Q10, 21 CFR 640 and Ph. Eur. 3.5. Multi-variate analysis and \$ quantification are produced from each manufacturer's own historical data and cost drivers, under a data-use agreement.