



SUPPLY CHAIN

NORTH AMERICAN MULTI-SITE MANUFACTURER

Cutting Stock at Network Scale

BACKGROUND

The client operates seven production facilities across Canada and the United States, serving a distributed customer base with overlapping capabilities across sites. By any conventional measure, the network was well-run. Each plant managed its own production efficiently. Orders were fulfilled. Customers were served. But the company's leadership suspected they were leaving significant margin on the table, but they could not pinpoint where.

The company's existing supply chain planning tools supported network-level decisions: which facilities should produce what, for which customers, at what cost. But those models treated production as interchangeable across sites. They could optimize for cost, capacity, and proximity. They could not account for the physical constraints that determine how raw material actually becomes finished product at each facility, or how much of it becomes waste.

That meant strategic recommendations about production allocation, customer assignments, and capacity planning were being made with no connection to what was executable on the shop floor. The company was committing to multi-million-dollar network decisions without knowing whether those plans were physically feasible, let alone optimal.

CHALLENGE

The cutting stock problem (how to cut raw material into ordered dimensions with minimal waste) is one of the most studied problems in operations research, with foundational work dating back to the early 1960s. At the individual site level, it is considered solved. Every plant either runs optimization software or has experienced operators who manage cutting patterns through years of accumulated judgment. Ask any plant manager and they will tell you: we know how to cut.

That confidence is earned. It is also the problem. When cutting stock is treated as a local, operational matter, the organization stops asking whether the right products are being cut at the right facilities in the first place. A plant can run perfect cutting patterns against the orders assigned to it and still be the wrong plant to fill those orders. The inefficiency is not in the pattern. It is in the assignment.

Operations leaders know their plants are not interchangeable: different machine widths, different cutter counts, different grade capabilities, different waste profiles on the same order. But the planning models they rely on for strategic decisions do not. The tools that decide which facilities serve which customers treat production as a uniform input. A customer assignment that looks optimal on a network map may generate significant waste on the actual production floor, because the model that made the assignment had no visibility into cutting feasibility.

The cost of this blind spot is larger than most organizations realize. Waste from trim loss is typically buried inside standard production costs, treated as a fixed overhead rather than a variable that changes with every cutting decision. Because it is not broken out, it is not managed. Suboptimal customer-to-plant assignments drive up waste, which drives up raw material consumption, which drives up cost-to-serve, which distorts the profitability picture that leadership uses to make the next round of network decisions. The errors feed forward.

The waste on the floor is not evidence of bad cutting. It is the physical residue of strategic decisions that were never validated against production reality.

At the network level, the problem becomes categorically harder. When seven factories with different machine configurations serve a shared customer base across two countries, the question shifts from “how do we cut efficiently at this mill?” to “which mill should make which product, for which customer, given the cutting constraints at every facility, and what is the most profitable answer across the entire system?” Prior approaches had addressed cutting stock across multiple plants at the production planning level. But no one had embedded cutting feasibility into the network optimization itself: the strategic layer where decisions about facility utilization, customer assignment, and distribution are made. That was the gap. And it was where the value was hiding.

SOLUTION

SimWell built a unified optimization model in anyLogistix that integrated cutting constraints directly into the supply chain network, connecting sourcing, production allocation, customer assignment, and distribution decisions with the physical reality of what each facility could actually produce. Cutting logic was not added as a downstream check. It was embedded as a core constraint within the network optimization itself, so that every recommendation the model produced had already been validated against real machine capabilities and real material physics.

The technical foundation required generating thousands of feasible cutting combinations per facility, based on each plant's machine specifications. Consider a single parent roll: depending on the machine it runs on, it can be cut into dozens of different width combinations, each yielding a mix of primary product, usable by-products, and a waste fraction. SimWell modeled every combination as a bill of materials, assigned full cost to the parent roll, and penalized waste as consumed raw material with no matching demand. For the first time, trim loss became visible to the network optimizer as a real cost that varies with every allocation decision.

This changed which answers the model produced. A plant might accept slightly higher waste on one order if doing so freed capacity for a higher-margin product elsewhere in the network. Another plant might take on orders it would never have received under proximity-based logic because its machine configuration made it the lowest-waste producer for that particular width. A third facility might be revealed as structurally unprofitable once cutting losses were accounted for. Information that proximity and capacity models had never surfaced.

It also changed how planning decisions were made. Before the model, production allocation was driven by capacity and geography, then validated informally against what operators knew they could run. With cutting embedded in the optimization, the planning team could evaluate scenarios (a facility shutdown, a shift in customer mix, a demand

surge in a specific product width) and know that every answer the model returned was executable. The conversation shifted from “can we do this?” to “which option is most profitable?”

The entire pipeline was designed for permanence. ERP data feeds automatically into the optimization engine, flows through post-processing, and surfaces in dashboards that production planners use directly. The model runs on a regular cadence, integrated with the client's forecasting and planning cycle. It does not require a dedicated analytics team to operate. As demand shifts, capacity changes, or customer portfolios evolve, the optimization updates with them.

RESULTS

By embedding cutting constraints into network-level optimization, SimWell gave the client something no planning tool had provided before: confidence that strategic plans were not just financially optimal on paper, but physically executable on the production floor.

Key project findings included:

- **Millions in annual profit improvement identified** through optimized production allocation, customer assignments, and network-wide waste reduction. These model-validated gains directly informed the client's planning decisions and capital commitments
- **5% reduction in material waste** by shifting trim loss from a fixed production cost to an active optimization variable managed across the full network
- **First integration of cutting constraints into supply chain network optimization**, embedding production-floor physics into the strategic layer where facility utilization, customer assignment, and distribution decisions are made, a gap that prior multi-plant approaches at the production planning level had not addressed
- **High-confidence scenario planning** for facility shutdowns, demand shifts, and customer rationalization, with cutting feasibility validated in every scenario before capital was committed
- **Repeatable digital planning process** with an automated ERP-to-model pipeline that runs on a regular cadence without dedicated analytics support

Based on the value the initial optimization delivered, the client extended the engagement, expanding the model to additional planning scenarios and incorporating it into their standard decision-making process for production allocation and customer strategy.

Millions

ANNUAL PROFIT
IMPROVEMENT IDENTIFIED

5%

REDUCTION IN
MATERIAL WASTE

7

FACILITIES OPTIMIZED
AS UNIFIED NETWORK

SimWell helps operators make better decisions under uncertainty. We build simulation and optimization models that reflect operational reality, test scenarios before committing resources, and transfer capability so teams can run decisions independently.