SOUR GAS OPERATOR

The bottleneck that wasn't

A sour gas operator was about to commit hundreds of millions to a terminal upgrade. A simulation found the real constraint two stations upstream, before the capex committee approved the wrong scope.

INTRODUCTION

Every operations leader eventually faces a decision where the obvious answer is the wrong one. The bottleneck everyone points to is a symptom. Hidden by interactions that no spreadsheet can resolve, the real constraint lies elsewhere in the system. Capital flows to the visible problem, but the throughput targets keep missing.

A sour gas operator faced one such decision. The question on the table was how to absorb a significant rise in upstream sulfur production. The path that looked obvious would have committed hundreds of millions to the wrong place in the system. The work was finding that out before the capex committee approved the wrong scope.

THE DECISION

The operator faced an aggressive corporate target. Production needed to rise significantly and stay elevated for the next 30 years. Sulfur came with that decision, whether the operator wanted it or not, as it does in any sour gas operation. More gas meant proportionally more sulfur to move, and if the sulfur could not move, the gas could not be produced. The sulfur itself carried commercial value as a chemical feedstock, but that was a secondary line. The case before leadership concerned protecting gas production.

The mechanism was straightforward. If the logistics chain backs up, the granulation plants run out of storage. When the granulation plants run out of storage, the gas plants stop producing. Pushing more gas through the upstream operation meant moving significantly more sulfur every day, through a logistics system already running close to capacity at current volumes.

The visible problem was the export terminal. Storage filled, ships waited at berth, and cleanup time kept consuming availability. Every operations team that walked through the facility reached the same conclusion. The terminal needed work, and the work was overdue. A terminal upgrade was the natural answer.

Rail capacity was the other open question. Operations ran two trains on a shared rail network operated by a separate company. Adding a third meant negotiating with the rail operator, sourcing additional crews, and absorbing operational complexity the team did not currently have to manage. Two threads were live in front of leadership: when to finally fix the terminal, and whether to commit to a third train. Neither conversation by itself reached the answer.

The capex implications ran into the hundreds of millions, with operating commitments stretching across rail, terminal, and shipping for the life of the asset. Leadership had to decide which combination of investments would deliver the target and which would absorb capital without solving the problem.

Spreadsheets could not answer that. Neither could intuition. The system had too many moving parts to hold in one head, and the interactions across those parts hid the real bottleneck behind the obvious one.

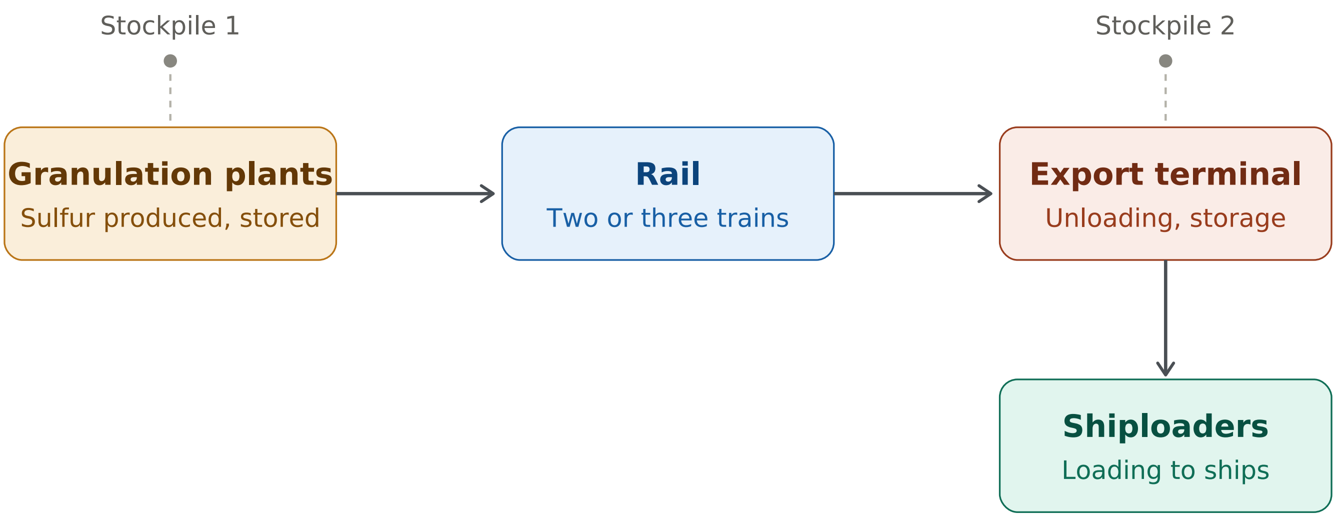

The sulfur logistics chain: production through rail to the export terminal, with stockpiles buffering at three sites.

WHY INTUITION WASN'T ENOUGH

The system ran on different clocks. None of them lined up.

Trains were the first clock. Their cycle times depended on loading speed at the granulation plants, unloading speed at the terminal, driver shift rules, and scheduled maintenance windows that the operator could influence but not control. Equipment failed on its own schedule. Stacker reclaimers went down for weeks at a time, conveyors tripped, and shiploaders broke. Each failure rippled differently depending on what else was happening in the system.

Ships were the other clock. A ship at berth that loaded slowly held the dock longer, delaying the next ship, which in turn delayed the one after that. Loading rates varied widely depending on the ship, the weather, and whether the previous ship's hatch changes ran short or long. The average loading rate looked like a single number on paper. The number hid a lot.

Stockpiles at three sites moved against each other. The unloading terminal filled when ships were slow and emptied when trains were slow. The granulation plants filled when the terminal was slow and emptied when trains caught up. A spreadsheet could compute averages, but it could not show what happened when a stacker reclaimer failed during a quarterly maintenance shutdown while a shiploader was offline for repair.

The operations teams knew the system intimately within their own functions. Rail operations knew the rail leg, the terminal team knew the terminal, and the shipping desk knew shipping. What no one had was a way to test the interactions across functions before reality tested them.

WHAT THE MODEL WAS BUILT TO ANSWER

Three questions framed the work.

Could two trains hit the lower throughput target? Operations ran two trains. If they could, the operator could defer the third-train decision. If they could not, the third train was a precondition for everything else.

Could three trains recover the system after an upset? When something went wrong, sulfur built up at the granulation plants while the chain worked through the disruption. After the upset cleared, the system had to run above normal rates to clear the backlog before another upset hit. A system that filled faster than it drew down drifted further behind every time anything went wrong. If three trains could not pull storage down fast enough, the operator needed to know what else had to change.

Which upset scenarios would force gas plant shutdowns, and how long would recovery take? The question that mattered most. The operator needed to know which combinations the system could ride through, which combinations required pre-positioned mitigation, and which combinations had no acceptable mitigation at all.

Each question demanded an answer that the operator could defend under the capex committee scrutiny.

THE FOURTEEN SCENARIOS

The operator's operations team came to the engagement with a list. Fourteen specific events they were worried could happen during normal operations. Each had been seen before in some form or had come close enough that it could not be dismissed. The list reflected operational anxiety from people who knew the system day to day.

The model tested each event against the three framing questions. Could the system absorb the event without forcing a gas plant shutdown? How long did recovery take? What changed about the answer if the event hit during normal operations versus during recovery from a previous upset?

The scenarios clustered into seven categories.

- **Planned maintenance on the unloading building.** Four days every quarter. Routine.
- **Stacker reclaimer failures.** Tested against a baseline where two of four units were already permanently out of service. Single-unit and multi-unit failures, varying durations.
- **Transfer tower outages.** Three towers, each carrying sulfur between unloading, storage, and the shiploaders. Each failure had a distinct downstream impact.
- **Unloading building shutdowns.** Seven-day and ten-day unplanned outages. A three-week scenario with one of two import streams down.
- **Shiploader failures.** From four days with both shiploaders unavailable to three months with a single unit down. The longest single-equipment outage on the list.
- **Export conveyor outages.** One and three weeks, modeled as equivalent to shiploader-stream failures.

- **Marine collision damage.** One berth out of service for a month.

The operator's team chose the list. The model tested every item on it.

WHAT THE MODEL SHOWED

Every scenario the operations team had put on the list went through the model. The results were sorted into four bands.

Survivable without intervention
Routine maintenance windows. No action beyond current practice.

Survivable but expensive
Single-equipment failures. Storage climbs, no shutdown. Drove overflow storage.

Forcing-function scenarios
Multi-equipment failures. System to the brink. Drove source-side fixes.

Unrecoverable under current configuration
Longest outages, worst infrastructure loss. Drove capital recommendations.

Severity increases top to bottom. The team chose the scenarios; the model sorted them.

Fourteen scenarios sorted into four severity bands, each band driving a different class of recommendation.

- **Survivable without intervention.** Routine maintenance windows sat in this band. The system absorbed the disruption in operating practice, and the model confirmed it. No recommendation flowed from these scenarios beyond what the team was already doing.
- **Survivable but expensive.** Single-equipment failures and shorter outages landed here. Storage at the granulation plants climbed during each event and recovery times stretched, but no scenario in this band forced a gas plant shutdown. The model quantified the remaining storage margin at each site for each scenario and identified the combinations that most narrowed the margin. The scenarios in this band fed the recommendation to add overflow storage at three sites.
- **Forcing-function scenarios.** Multi-equipment failures and longer outage durations drove the system to the brink. Storage filled at one or more sites before the system could recover. The model put a value on time-to-shutdown for each. Fixing the symptoms at the terminal did not meaningfully extend that number. The scenarios in this band drove the source-side recommendations: address dust at the granulation plants, replace equipment that the duty had outgrown, and plan for three-train operation as the baseline.
- **Unrecoverable under the current configuration.** The longest single-equipment outages and the worst infrastructure-loss events had no acceptable mitigation without infrastructure investment. The model carried the system through each scenario to its breaking point and reported back the conditions that forced the break. The scenarios in this band drove the capital recommendations.

Three findings cut across the bands.

Two trains could not deliver the lower target.

The model ran two-train operations against every reasonable combination of rail cycle time, equipment availability, and operational discipline. The system capped well below the lower target in every run. The math was the constraint. Two trains running a full cycle of loading, transit, unloading, maintenance, and driver swaps could not complete often enough to deliver the required daily tonnage. Pushing the cycle time down to physically achievable limits still came up short.

Adding a third train became a precondition for any target above current operations.

Three trains could hit the higher target, but only at the upper edge of feasibility. Recovery was unreachable.

Three-train operations met the higher-throughput target in the model, but the system ran hot doing it. Berth occupancy climbed into a range that left little room for shipping upsets the model did not test for, like weather delays and ship breakdowns. The unloading building operated close to continuous availability. Storage at the granulation plants sat higher than preferred, leaving a thinner margin to absorb the next upset.

Recovery was a different conversation. To clear storage that had built up at the granulation plants during an upset, the system needed to run well above the higher throughput target for a sustained period. The model capped well below that rate, even with three trains at the upper edge of feasibility. Hitting it would require a system redesign. The current configuration would not get there at any utilization.

The bottleneck originated two stations upstream.

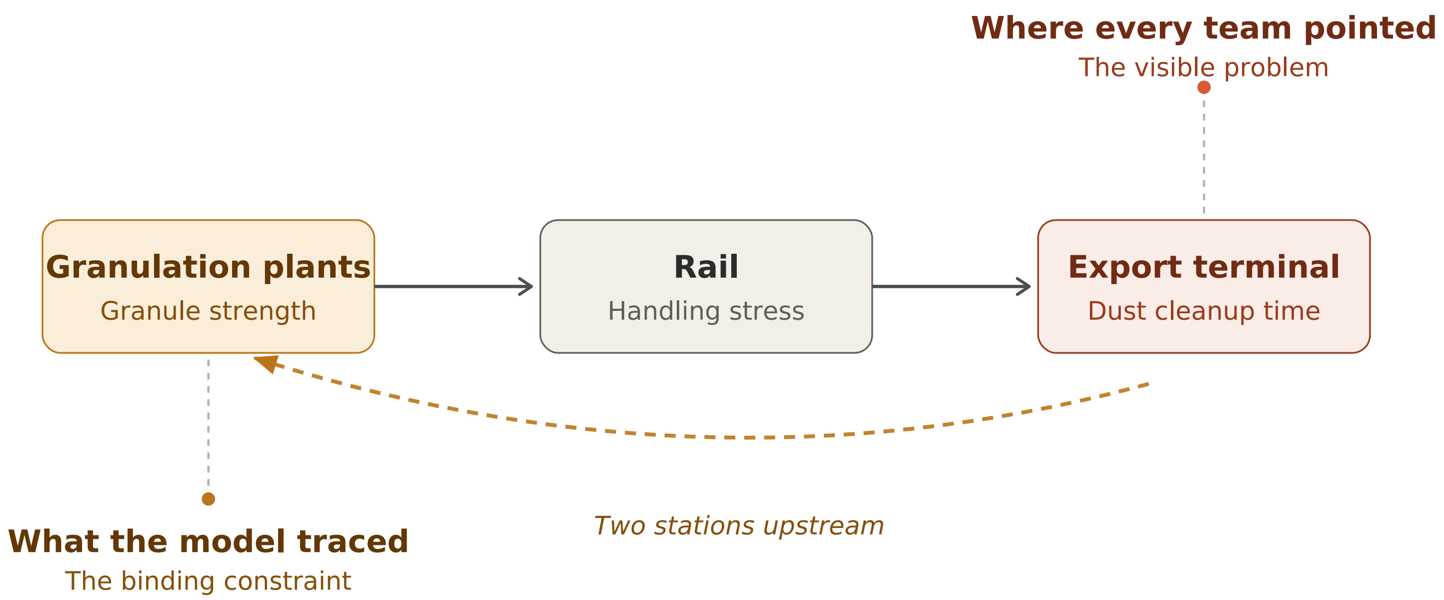

The visible problem sat at the terminal. The binding constraint sat two stations upstream, at granule strength.

The operations teams knew the terminal cleanup time was excessive. They had been living with it. What the model traced was the source of the cleanup time.

The unloading building at the terminal was the visible problem. It went offline regularly for cleanup and ran below design rates. Its equipment failed more often than the original specifications suggested it should. The daily evidence pointed clearly at the terminal, and the teams that worked there had been pointing at it for years.

The model traced the failure pattern backwards. The equipment along the chain could not process sulfur at the volumes required without generating excessive dust. Cleaning the dust at the terminal consumed roughly a quarter of the system's availability. The dust was the residue of granules that had broken down across the handling chain. Loading, transit, and unloading each contributed mechanical stress, and granules whose strength fell below the engineered range did not survive. The granulation plants were producing those granules. The terminal's dust extraction had been sized for product specifications that the chain was not delivering.

The granulation team had no signal at production that anything was wrong. Their product met production tests. The terminal team saw the fines arrive and reported the terminal as failing. Neither team had visibility into what the model showed: granule strength at production was the variable controlling cleanup time at the terminal.

The same pattern held for the other major reliability constraints. Stacker reclaimers, transfer towers, conveyors. Each had been specified for an isolated function. Each was now operating in an environment its original design had not anticipated, shaped by upstream and downstream conditions the system was never tested against.

The simulation produced comparative indicators. The values were not absolute. A reader looking for a precise throughput prediction from any single run would be reading the work wrong. The findings that mattered were directional and structural: which scenarios the system could absorb and which it could not, and where the binding constraints lay. The operator's team understood the distinction and used the findings as intended by the model.

A terminal upgrade alone would not deliver the throughput targets. The fix had to span four operating systems: granulation plant product quality, rail cycle time, terminal reliability, and shiploading throughput. Each was binding. None was sufficient on its own.

THE RECOMMENDATIONS

The recommendations followed the model's structure. Four operating systems, each addressed on its own terms, each tied to a specific decision the operator had to make.

Address dust at the source.

The terminal's dust management systems had been sized for a smaller load than the granulation plants were sending downstream. Cleanup after unloading consumed roughly a quarter of the terminal's availability on its own. Product quality at the loading point was the binding constraint.

Evaluate granulation plant output for dust and lumps before the rail cars leave the loading yard, and invest in screening, conditioning, and handling to reduce dust generation at the source. Screening would also remove the lumps that slowed unloading at the receiving end. The downstream dust collection systems at the terminal still needed replacement, including baffles in the dump grid and a properly sized dust extraction system, but as a second order of work.

Replace, don't repair.

Roughly twenty equipment items across the system showed availability well below what their original specifications promised. The specifications did not match the duties they were being asked to do. Dust systems at the unloading building had been designed for a different load profile. Equipment installation choices made the systems harder to maintain than they needed to be. The pattern was consistent across the list.

The recommendation was to engage a competent engineering review of each item, identify whether the failure mode reflected wear, design, or environment, and replace or redesign accordingly. Like-for-like replacement would reproduce the same failures. Operator training on the new configurations was scoped alongside the equipment work.

Add overflow storage at three sites.

Even with the four-system fix, certain upsets would still drive storage levels above current capacity for short periods. The model showed which upsets, at which sites, and for how long. The recommendation was to add temporary granule storage at both granulation plants and the terminal, sized to accommodate the worst-case scenarios identified by the model. Granule storage was cheaper than blocking sulfur or building new permanent infrastructure, and it bought the system time to ride through events that would otherwise force gas plant shutdowns.

Reduce ship turnaround.

Ship turnaround time had already improved from earlier baselines. The operations teams had worked it down over time. The model showed further reductions were achievable and would relieve pressure on the export end of the system. Faster continuous loading rates, fewer non-loading hours at berth, and better coordination on hatch changes and draft surveys could move turnaround into a range that lowered berth occupancy from stressed to sustainable. Each hour out was an hour the system could spend drawing down storage after the next upset.

Plan for three-train operation as the baseline.

Three trains were a precondition for the throughput targets. The operator's existing planning treated three trains as a recovery mode, used occasionally to draw down storage after upsets. The recommendation was to invert that and make three trains the normal operating pattern. The work of negotiating with the rail operator, sourcing additional drivers, and rebuilding the scheduling framework had to start immediately, because nothing else in the recommendation set would deliver the targets without it.

The recommendations arrived as a sequenced capex plan with each item tied to a specific operational failure mode the model had quantified. The capex committee had what it needed to commit capital where it would move the throughput, and to defer or reject the items that would not.

RESULTS

- Capex committee informed across hundreds of millions in committed capital.
- Bottleneck traced two stations upstream from where every operations team had been pointing, surfacing a binding constraint no single function had visibility into.
- Three-train operation identified as a precondition for any throughput target above current operations.
- Fourteen operational upset scenarios sorted into four severity bands, with time-to-shutdown values and storage-margin quantification on the consequential ones.
- Recommendation set spanning four operating systems, sequenced as a capex plan the committee could approve, defer, or reject item by item.

WHAT FOLLOWED

In the period after the study, the operator made the decisions that the model had been built to inform.

The upstream expansion that the study had stress-tested reached the final investment decision. The production wave became committed capital. Three-train operation became the operator's planning baseline.

The operator tendered a capacity expansion at the export terminal. The published scope addressed the operating systems the study had identified as binding constraints: dust collection, unloading building modifications, drainage, and rail integration.

The tender documentation the operator published for the project named the upstream cause directly. Reliability problems at the terminal stemmed from product handling and from dust systems that the volumes had outgrown. The operator's diagnosis matched the model's. The terminal scope addressed a symptom that the operator now described in the same terms the simulation had traced.

The chain is moving the volumes the operator committed to.

WHAT THIS KIND OF WORK DELIVERS

The operator had committed to a production target that the existing chain was not built to carry, sustained over thirty years, with the gas business depending on whether sulfur could move at the new volumes. The instinct was a terminal upgrade. The visible problem looked solvable that way, and the capex committee was prepared to fund the work.

The model showed that the terminal upgrade alone would not have delivered the target. Hitting the throughput goal required investment across four operating systems, and no single one of them would have been sufficient on its own. The terminal scope absorbed capital that would not have moved the targets without upstream changes. The upstream changes were the work that no team had been requesting, because no team had the visibility to see what the chain needed as a whole. The model showed the combination.

The capex conversation moved from “should we upgrade X” to “which combination of upgrades across X, Y, Z, and W actually delivers the throughput, and which combinations protect upstream production when something fails.” The model put defensible numbers in front of the decision the committee needed to make.

SimWell helps oil and gas operators see the system as a whole, before reality tests it. We build simulation and optimization models grounded in real operating reality, then package the work into decisions leadership can defend.