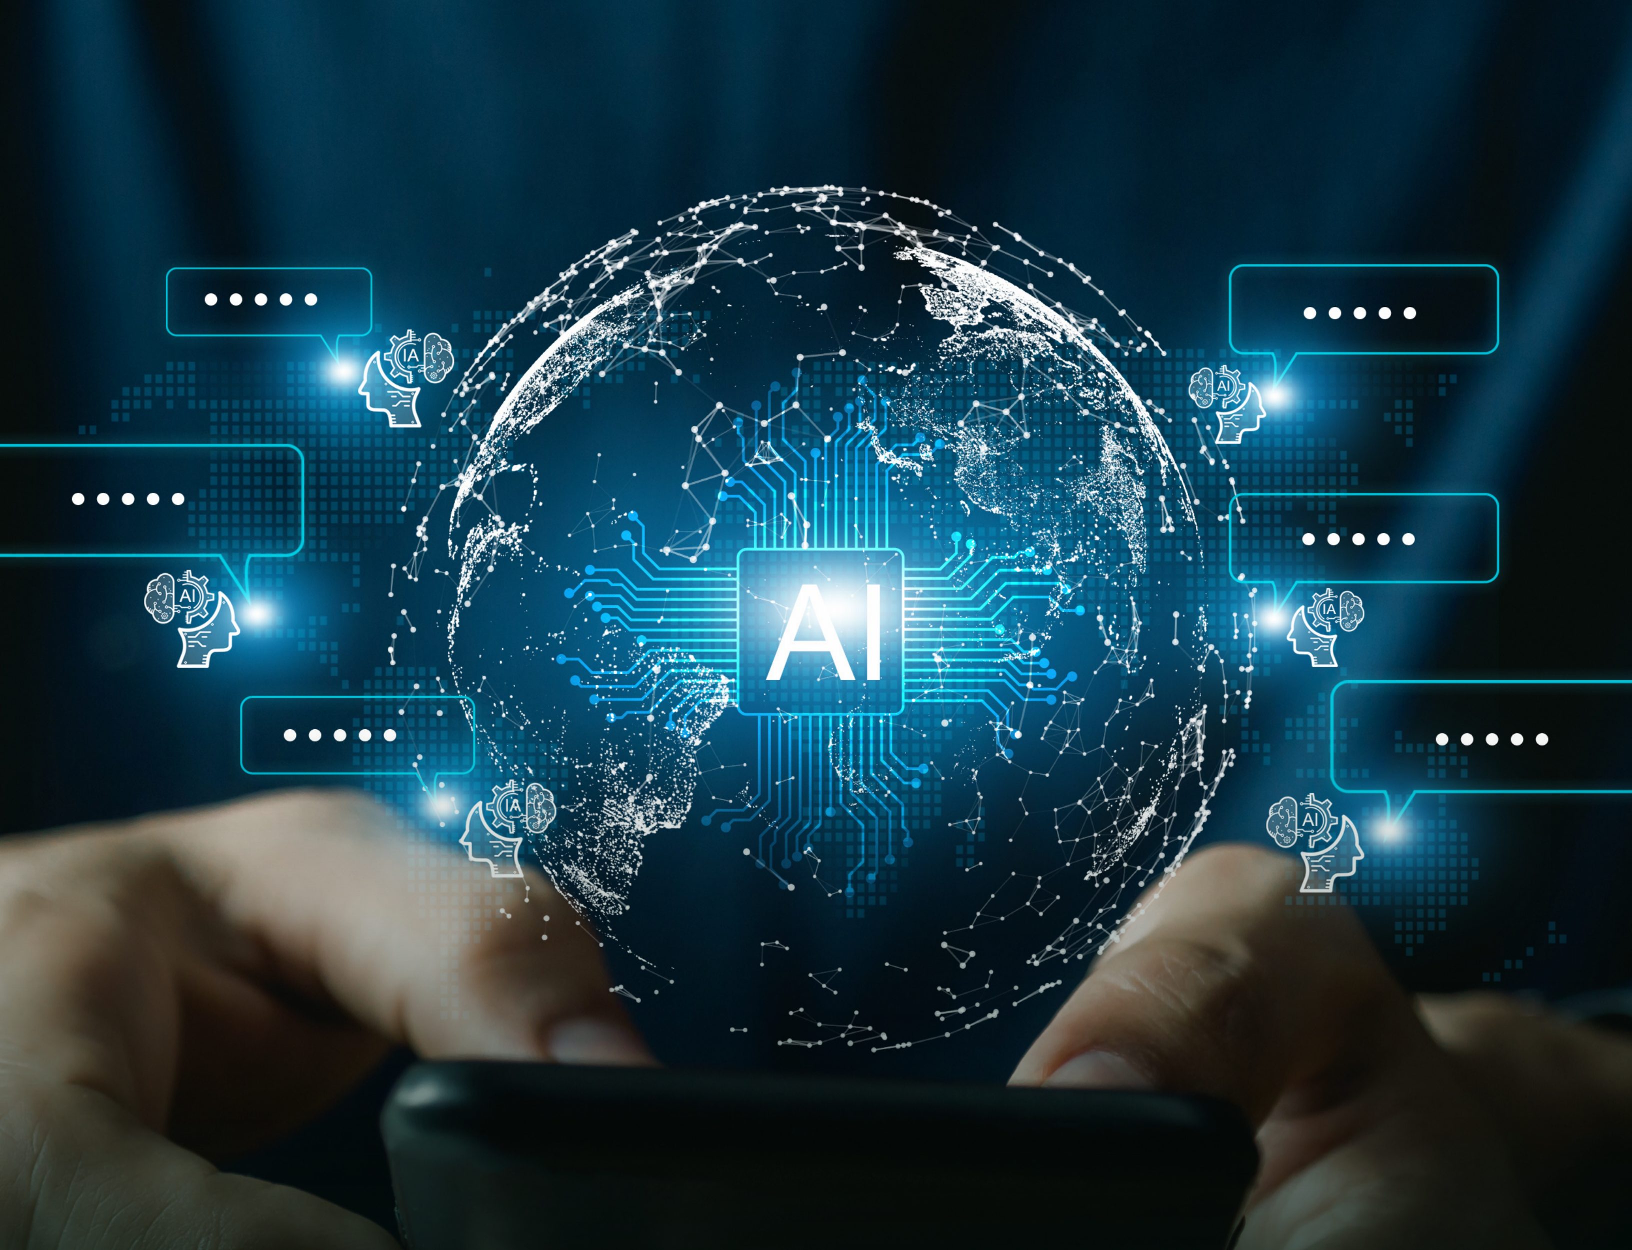


GAINS
MOVE FORWARD FASTER

Whitepaper

# Less Talk, More Execution

## Turning AI Into a Decision Engine, not just a chat bot.

Whitepaper

# How GAINS policy-bound AI improves decision speed without surrendering control

## The Old Rules Are Gone. Are You Still Playing By Them?

Here's something worth sitting with for a moment: supply chain planning worked well enough for a long time. Not because the people in it were brilliant, though many were, but because of the (mostly) stable operating environment. Demand moved in recognizable patterns. Lead times were long, but they were predictably long. Suppliers, carriers, and customers formed a kind of shock-absorbing ecosystem. When something went wrong, the system had enough slack to recover before it caused too many issues.

### That era is over, and it didn't end well.

The pandemic didn't create supply chain fragility; it revealed how fragile it has always been, compressing years of latent disruption into a few chaotic months and taking years to recover. What has come after isn't a return to normal, but the arrival of a new permanent baseline: lead times that swing by a margin of weeks without warning, demand signals that go stale within a planning cycle, geopolitical shifts that can restructure a sourcing network overnight, and network complexity that grows faster than anyone can manage with spreadsheets and a planner's manual heroics.

Today, supply chain executives don't need more reports or better dashboards. What they need is a system that can think fast enough to keep pace, that absorbs volatility, makes better decisions faster, and keeps the enterprise aligned when conditions change faster than any planning cycle can respond. That's the true problem AI needs to solve.

**The uncomfortable truth is most AI being sold into the supply chain market today doesn't come close.**

Find Out Why It Doesn't Deliver

## Why The Current AI for Supply Chain Playbook Can't Deliver

AI technology continues to move quickly, and the AI vendors claim they can make moves even faster. If you've been to a supply chain conference recently, you've noticed that everyone has "AI" now. But on closer inspection, most of what's being sold falls into one of two familiar categories.

- **AI Insights** — dashboards, anomaly alerts, natural language interfaces, chat-based assistants that surface information and answer questions. Genuinely useful tools, however, structurally, incapable of making decisions. Knowing there's a problem and knowing what to do about it are different tasks entirely.
- **Point prediction AI** — machine learning applied to demand or lead time forecasting, delivered as a standalone output, disconnected from the decision logic that should consume it. The thing about a better forecast is that it doesn't change what gets ordered or when: it's not an improvement in decision making. It's an analytical exercise with better graphics.

Both have value. Neither one addresses the core problem.

The difficult part of supply chain planning has never been generating a forecast or flagging anomalies. The hard part has always been deciding what to do when

- constraints collide, and service targets conflict with working capital limits
- supplier shortages require network-wide rebalancing across hundreds of SKUs and dozens of locations simultaneously
- when a short-term expedite creates a downstream cost problem that won't surface for weeks

Supply Chain decisions like these require reasoning across constraints and explicitly quantifiable trade-offs. The system must connect recommendations to execution at speed and scale, and keep a human in the loop who is accountable for the outcome. A chat interface can't do that. A standalone forecast can't do that. A dashboard can't do that.

What supply chain executives need is a **decision system** where AI isn't a feature update, but an integrated capability embedded in how decisions get made, validated, executed, and improved over time.

## Non-Negotiables of Modern Supply Chain Decisioning

Before evaluating any AI tool, it's worth being precise about what a supply chain decision system needs to accomplish. Because precision here is what separates a genuine capability from a well-marketed feature. It must predict uncertainty accurately enough to act on.

Forecasts, lead times, and supply variability need to be estimated with enough precision and transparency that downstream systems can consume them automatically, and planners can tell when to trust recommendations and when to investigate. **A prediction that isn't calibrated for action is just analysis.**

## An Effective Supply Chain Decision Engine Must

### Optimize trade-offs under real constraints

The best supply chain decision is the one that balances service levels, inventory investment, capacity utilization, and working capital simultaneously, under the actual constraints of the network. That requires embedded optimization, not just prediction. And the combination matters: prediction without optimization produces recommendations that might be the best fit locally but systemically damaging.

### Automate where policy is clear

Repeatable, routine decisions (replenishment orders within defined parameters, safety stock updates, standard allocation logic) should not require human review. They should execute automatically, within policy-defined guardrails, and route only genuine exceptions to planners. A system that sends every decision to a human for approval isn't a decision system. It's expensive noise.

### Orchestrate across time horizons and functions

A replenishment decision made today that quietly creates a capacity constraint next quarter isn't a good decision — it's a locally optimal one. Near-term execution must align with mid-term planning and long-term network design. Cross-horizon orchestration is what prevents the short-term optimization trap.

### Govern every decision

Every recommendation, every automated action, every exception needs to produce an auditable record: what was decided, when, why, what assumptions were used, and what actually happened. Without a decision log, there's no learning loop, no accountability, and no mechanism to get better over time. Governance isn't bureaucracy; it's the infrastructure that makes trust possible.

### Keep humans in control where judgment is required

**That doesn't mean humans review everything!** — That's not control, that's a bottleneck. It means the system knows precisely when to act, when to escalate, and when to explain, and it gives planners the context they need to exercise judgment quickly and with confidence.

## Practical AI, Built for Decisions: Purposeful, Embedded, Accountable

GAINS believes in a simple idea: **AI is a tool, not a strategy.**

The question supply chain leaders must ask themselves is. "What supply chain decision needs to improve, and what technique best reduces the uncertainty or complexity involved?"  Not, "Where can we put AI?" Our philosophy produces a deliberate, layered pragmatic approach to these areas:

### Where uncertainty lives

**In demand signals, lead time variability, and supplier behavior**: machine learning improves prediction accuracy and replaces static, manually maintained parameters with dynamic values that update continuously as new data arrives. Because the world doesn't stay still, and neither should your planning inputs.

### Where trade-offs live

**Across service, cost, capacity, and inventory**: operations research and optimization make trade-offs explicit, feasible, and aligned with actual business policy. This is the layer most AI vendors skip, but it is the most important. ML-based prediction without Operations Research or, (OR-based optimization) produces recommendations that are potentially brilliant in one dimension while quietly catastrophic in others.

### Where decisions are repeatable, and policy is clear

**Replenishment, parameter updates, order execution**: The system automates routine decisions within defined guardrails, achieving high auto-approval rates while routing true exceptions to planners with context. The planner's job becomes review, not triage.

### Where alignment across functions and horizons matters

Decisions are orchestrated so that near-term execution connects to mid-term planning, and a decision in one corner of the network doesn't create a problem three time zones and two quarters away.

### Where human judgment is genuinely required

Exceptions arrive with trade-offs already quantified, constraints already visible, and recommended actions already rationalized. Planners make better decisions faster because the analytical work is already done when they open the screen. The result is not a collection of modules, but a solution stack consisting of prediction, optimization, automation, and orchestration functioning as a coherent system.

> "What supply chain decision needs to improve, and what technique best reduces the uncertainty or complexity involved?" Not, "Where can we put AI?"

## DEO Isn't a Feature — It's an Architecture

Here's something that people pitching AI solutions don't always say: any vendor can sell you an AI feature. Building a decision system that makes AI safe, scalable, and operationally valuable is a fundamentally different problem, and it requires assets that take decades to develop.

**Embedded optimization engines.** GAINS has deep operations research capability built directly into its planning engine. This is the mechanism that makes trade-offs explicit and feasible within actual network constraints. Without it, AI recommendations are unconstrained — potentially optimal on the dimension you're measuring, potentially damaging to the dimensions you're not.

**Historical transactional and execution data.** GAINS learns from what actually happens — not just what was planned, but what was ordered, received, shipped, and held. Execution data is the foundation for models that improve continuously and for learning loops that connect outcomes back to the assumptions that drove them. That feedback loop is how systems get smarter rather than just faster.

**Policy frameworks and guardrails.** Every automated decision in GAINS operates within a defined policy structure. Guardrails establish the boundaries within which the system acts autonomously. Outside those boundaries, decisions route to humans. This architecture is what makes automation genuinely safe — not safe in theory, but safe in production, at scale, with real consequences.

**Decision logs and auditability.** GAINS maintains a complete record of what was decided, when, why, and with what outcome. This is what enables continuous improvement, supports regulatory and operational accountability, and gives supply chain leaders actual visibility into how their network is being run — not how it was planned to be run.

> AI that sits outside the system of record generates recommendations that get ignored.

**Cross-horizon orchestration.** The alignment of tactical and strategic goals across time horizons and functions prevents the compounding inefficiencies that emerge when decisions are optimized locally but not globally a problem that's almost invisible until it's expensive.

**Domain models built for supply chain reality.** GAINS models are purpose-built for inventory positioning, replenishment, multi-echelon trade-offs, capacity allocation, and service level management. Generic AI applied to supply chain data is not the same thing as AI designed around supply chain decision logic. The domain knowledge is load-bearing.

**Workflow integration.** GAINS is embedded in the planning and execution workflows that customers already operate. AI that sits outside the system of record generates recommendations that get ignored. AI integrated into the system of record drives adoption, accountability, and outcomes that show up in financial statements.

## Where GAINS Applies AI Today

### 1. Lead Time Prediction

Static ERP lead times are one of the most quietly consequential sources of planning error in complex networks. They get set once, updated infrequently, and in most organizations bear little resemblance to what suppliers are actually delivering. They're a form of institutionalized optimism.

**GAINS replaces static lead times with AI-estimated values**, trained on historical supplier performance, additional data outside data (vendor attributes, item attributes, external variables, etc.), and updated continuously as new execution data arrives. When confidence thresholds are met, updated values are applied automatically inside the planning system of record, no manual intervention, no lag. When variability is elevated, those signals propagate directly into safety stock calculations and order timing decisions.

Implementation matters too: GAINS' Lead Time Prediction tool was delivered to Border States as a plug-in microservice that integrated directly with SAP — no ERP overhaul, no heavy IT engagement, time-to-value accelerated by nearly a full year.

**Outcomes**: Border States, one of the top 10 electrical distributors in the U.S., manages 750,000 SKUs across 130 warehouses and a $600M inventory network. Before GAINS, lead times were tracked in a homegrown Access database — a manual process that created blind spots, stockouts, and excess inventory simultaneously. After: a 31% reduction in lead time error, a 65% improvement in lead time accuracy, and a 30% decrease in expediting costs as better visibility eliminated a significant share of costly last-minute rush orders. Knowing when your

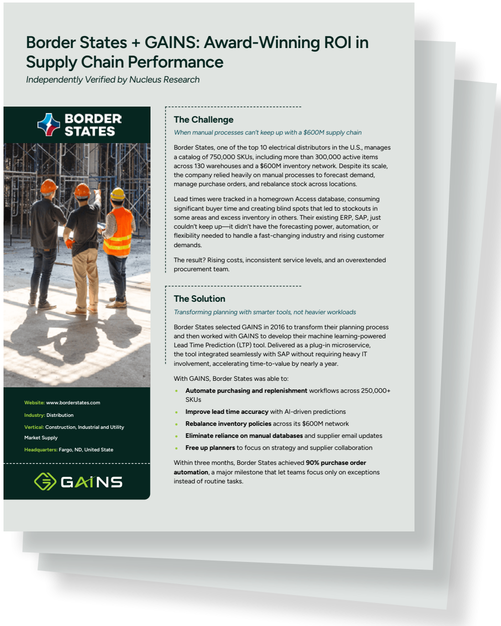

Border States + GAINS: Award-Winning ROI in Supply Chain Performance

*Independently Verified by Nucleus Research*

BORDER STATES

**The Challenge**

*When manual processes can't keep up with a $600M supply chain*

Border States, one of the top 10 electrical distributors in the U.S., manages a catalog of 750,000 SKUs, including more than 300,000 active items across 130 warehouses and a $600M inventory network. Despite its scale, the company relied heavily on manual processes to forecast demand, manage purchase orders, and rebalance stock across locations.

Lead times were tracked in a homegrown Access database, consuming significant buyer time and creating blind spots that led to stockouts in some areas and excess inventory in others. Their existing ERP, SAP, just couldn't keep up—it didn't have the forecasting power, automation, or flexibility needed to handle a fast-changing industry and rising customer demands.

The result? Rising costs, inconsistent service levels, and an overextended procurement team.

**The Solution**

*Transforming planning with smarter tools, not heavier workloads*

Border States selected GAINS in 2016 to transform their planning process and then worked with GAINS to develop their machine learning-powered Lead Time Prediction (LTP) tool. Delivered as a plug-in microservice, the tool integrated seamlessly with SAP without requiring heavy IT involvement, accelerating time-to-value by nearly a year.

With GAINS, Border States was able to:

- **Automate purchasing and replenishment** workflows across 250,000+ SKUs
- **Improve lead time accuracy** with AI-driven predictions
- **Rebalance inventory policies** across its $600M network
- **Eliminate reliance on manual databases** and supplier email updates
- **Free up planners** to focus on strategy and supplier collaboration

Within three months, Border States achieved **90% purchase order automation**, a major milestone that let teams focus only on exceptions instead of routine tasks.

Website: www.borderstates.com

Industry: Distribution

Vertical: Construction, Industrial and Utility Market Supply

Headquarters: Fargo, ND, United State

GAINS

suppliers will actually deliver, it turns out, is worth a great deal, receiving 976% ROI on its LTP solution with payback in less than six weeks.

### 2. Demand Prediction

GAINS applies feature-engineered machine learning to demand forecasting at the SKU-location level (SKUL), across multiple time horizons. Models incorporate promotion uplift, cannibalization effects, halo modeling, and anomaly detection as well as the incorporation of data from internal and external systems including macro factors and consumer behavior. Critically, they include calibrated confidence signals, so planners understand not just what the forecast says, but how much to trust it and when to investigate the signal behind it.

Forecasts aren't delivered as standalone outputs requiring a manual handoff to a separate planning system. They're natively integrated into the planning system, automatically influencing downstream inventory and supply decisions. The forecast doesn't just inform the decision; it's part of the decision.

## 3. Supply Decision Automation

A mid-size distributor makes tens of thousands of buying decisions every week. A large manufacturer operating across multiple product lines and locations makes more. No planning team, regardless of size, experience, or tooling, can meaningfully review decisions at that volume. The math doesn't work. What happens instead is that planners spend their time on the decisions in front of them; routine orders get approved based on instinct or habit, and the decisions that actually warrant scrutiny get the same attention as the ones that don't.

That is the problem GAINS Supply Decision Automation solves — not by following rules, but by applying machine learning to the full complexity of each buying decision: evaluating supplier options, inventory transfer alternatives, order quantities, service commitments, and cost constraints simultaneously, and executing the right decision automatically at scale.

Purchase orders, transfer orders, and work orders are generated, validated against guardrails, and executed without requiring a planner to review every line. Only genuine exceptions — decisions where the system's confidence is low, constraints conflict, or thresholds are breached — are routed to planners, with trade-offs already quantified and context already assembled. The result is not a team working harder to keep up with decision volume. It is a system that handles the volume, and a team focused on the decisions that actually require them.

## 4. Agentic Decision Orchestration

GAINS' agentic capabilities represent the deliberate next step in governed decision engineering. When we use the term agentic, we don't refer to autonomous agents operating outside defined constraints, making independent decisions and hoping for the best. The agentic AI used as part of GAINS DEO are policy-bound, constraint-aware agents designed to monitor conditions, simulate downstream impacts, propose coordinated actions, and route decisions through defined approval workflows, all within GAINS' optimization and governance architecture.

**In practice**, A supplier delay signal is detected. The agent simulates downstream inventory impact across affected SKUs and locations. It evaluates rebalancing options against service targets, cost constraints, and available capacity. It surfaces a recommended action with trade-offs quantified. If the action falls within the defined policy, it executes. If it requires human review, it routes to the appropriate planner, **not just an alert, but a decision-ready recommendation with full context assembled.**

## AI Chat Agent Vs. Decision Agent

The distinction between a chat assistant and a decision agent should be clear. A chat assistant lives in the conversation layer: it explains, answers questions, and helps users navigate the system. A decision agent lives in the decision layer. It owns a bounded responsibility, maintains state across time, acts on events rather than waiting to be prompted, and produces real operational consequences.

A system that explains infeasibility is useful. A system that owns feasibility as a responsibility (detecting it, diagnosing root cause, proposing and applying corrections, managing the resolution loop changes outcomes. The supply chain leaders must build a tight integration between AI capability and decision governance.

## Build vs. Buy: Why the Hard Part Isn't the Model

Every supply chain organization with a data science team is asking the same question right now: "Shouldn't we just build AI capabilities internally, or extend the platforms we already have?"

The honest answer is that building machine learning models has never been more accessible. Data science teams can build demand forecasting models, lead-time prediction models, supplier risk classifiers, and feature pipelines using modern tools with reasonable effort. That part of the AI stack in the prediction layer is no longer a hard problem. It is achievable, and in many cases, teams are already doing it.

The hard part is what comes next.

**Prediction signals don't run a supply chain**. A predicted demand value, a predicted lead time, a supplier risk score — these are outputs. They have no operational value until they are connected to the system that is actually making supply chain decisions: the planning policies, optimization engines, workflow orchestration, and execution logic that determine what gets ordered, from where, in what quantity, and when.

**Most internal AI initiatives stop at the prediction layer**. They successfully produce signals. What they cannot easily do is integrate those signals into an operational decision system in a way that changes outcomes because that integration is not a data science problem. **It is an architecture problem**, and it requires assets that take years to build.

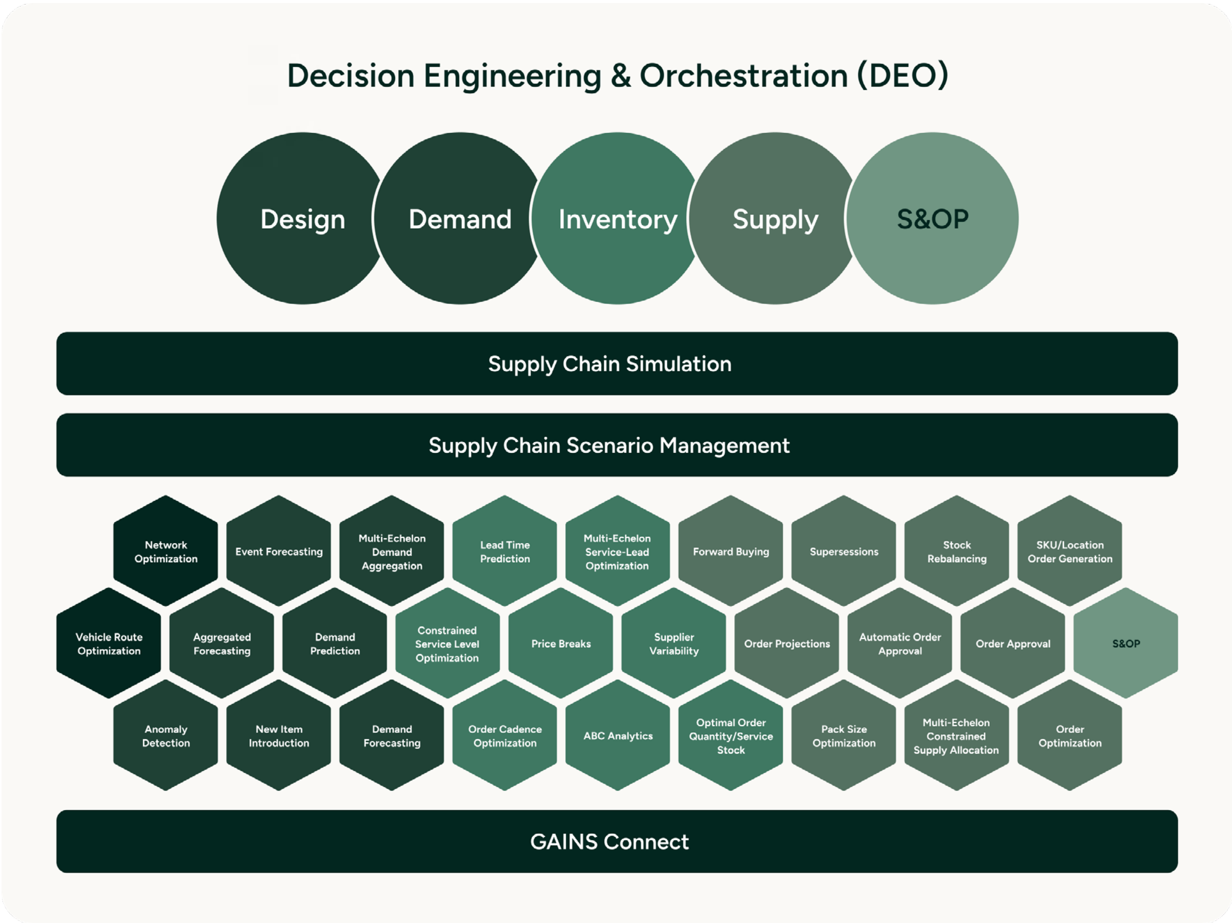


## Two examples illustrate why this gap matters in practice.

**Lead Time Prediction**. Predicting supplier lead times with machine learning is valuable on its own. **But the real value only appears when those predictions feed directly into supply planning decisions**, changing safety stock positioning, adjusting order timing, and influencing supply feasibility calculations. As planning moves toward probabilistic models, the output from lead time prediction becomes a direct input into the planning engine itself. That requires tight integration between the ML model and the planning logic. A prediction model that lives outside the planning system of record cannot operationalize those improvements. The signal exists. The decision doesn't change.

**Supplier Decision Automation**. GAINS uses machine learning to increase the first-pass approval rate of purchase orders generated by the planning system. But in developing that capability, what became clear is that **the model was only part of the solution**. To make automated recommendations reliable, the decision logic inside the planning system itself had to be refined; specifically, the inventory sharing and transfer logic had to be brought into the order optimization process. That refinement ensures the system evaluates the full range of alternatives, supplier sourcing options, inventory transfers, and order quantities — before approving a purchase order automatically. That level of integration is not possible if the ML model sits outside the decision system. **The model can recommend, but only an integrated system can decide.**

This is the distinction between prediction and decision intelligence. Prediction produces signals. Decision intelligence uses those signals to change what the supply chain actually does within policy, under constraints, with governance, and connected to execution.

Building the prediction layer internally is achievable. Building the integrated decision system that makes predictions operationally valuable is not, certainly not quickly, and not without the domain models, optimization engines, policy frameworks, and execution integration that GAINS has spent years developing.

The build vs. buy question, properly framed, is not "can we build ML models?" It is "Can we build the decision architecture that makes those models matter?" For most organizations, that answer changes the calculus entirely.

## Where DEO Delivers the Biggest Payoff

GAINS Pragmatic AI delivers the highest value where decision complexity and volatility are most acute and where the cost of a bad decision at scale is most visible on the balance sheet.

- **Inventory-intensive distributors and manufacturers** with high SKU counts and service commitments that cannot tolerate planning errors at scale. When you're managing 250,000 active SKUs across 130 warehouses, the math on manual processes doesn't work.
- **Multi-echelon networks**, where allocation and deployment trade-offs cascade across locations and create systemic risk when managed manually. What looks like a local inventory problem is often a network design problem in disguise.
- **High-volatility environments** (lead time swings, demand unpredictability, supply shortages) where static parameters and manual overrides are structurally insufficient. You can't solve a dynamic problem with a static answer.
- **Acquisition-driven network complexity** where inconsistent policies, fragmented data, and misaligned planning practices compound into systemic inefficiency. As illustrated by:

**Continental Battery**: 150+ locations across N. America, each with its own planning logic, required not just better forecasting but a unified decision architecture and a network design capability that could model facility roles, sourcing strategies, and inventory placement simultaneously.

**The results**: 40% inventory reduction, fill rates rising from 65–70% to the high 80s and low 90s, a projected 10% operational cost reduction through asset optimization, and a 10–13% working capital reduction. without adding new facilities.

- **Organizations with entrenched override culture**, where planner heroics mask underlying decision quality problems rather than resolving them. When the system can't be trusted, people stop trusting it, the overrides multiply, and then you can't tell whether the system is wrong or whether the overrides made it wrong.  Breaking that cycle requires a decision system people actually trust.

## Outcomes Leaders Can Count On

GAINS turns planning into a faster, more reliable operating system for supply.

- **Higher service levels with fewer expedites**. Decisions are more accurate, more consistent, and more tightly connected to execution, which reduces the surprises that force costly last-minute interventions. Fewer fires mean fewer fire drills.

- **Lower working capital**. Better prediction and embedded optimization reduce the "just in case" inventory that accumulates when planners don't trust their systems. Buffers get sized to actual variability, not worst-case assumptions. That's not a small number again. Border States freed up $21M in six to eight months.

- **Greater operational stability**. Fewer plan swings. Fewer crises. Faster recovery when disruption hits — because the system can re-optimize quickly rather than waiting for a manual replanning cycle that takes days.

- **Planner productivity reallocation**. Less time on routine decisions that a well-governed system should be making automatically. More time on the exceptions that genuinely require judgment, the supplier relationships that require human relationship skills, the strategic questions that require experience. Planners become decision managers rather than order processors.

- **Decision velocity and auditability**. Faster planning cycles, with transparent trade-offs and governance structures that give leaders confidence in the decisions their teams are making and the audit trail to verify outcomes over time. That confidence, it turns out, is its own form of competitive advantage.

## From Recommendations to Resolution: The Next Phase of DEO

Prediction without optimization produces unconstrained recommendations. Automation without governance produces operational risk. Agents free of decision architecture produce noise.

The future of AI in the supply chain isn't about how much autonomy you can give an AI system. It's about building a decision architecture that makes AI genuinely useful and safe, in the face of increasing levels of complexity and speed. **GAINS is deliberately building toward that standard, guided by the principle of credibility before autonomy.**

**Decision readiness agents** that assess whether data, parameters, and network inputs are an actual fit for planning use scoring confidence, surfacing structural risks, and flagging conditions before they propagate downstream errors. Think of these as the agents that ask, "are we ready to make this decision?" Before the planning run starts.

**Infeasibility and exception agents** that don't just detect problems but own the resolution loop, diagnosing root cause, simulating corrective options, applying fixes within policy, and escalating only what genuinely requires human judgment. The difference between flagging a problem and resolving one is the difference between a notification and a decision.

**Scenario and trade-off agents** that frame decision alternatives, quantify impacts across service, cost, and capital dimensions, and present structured recommendations rather than raw outputs. Less "here's what the data says," and more "here's what your options are, and here's what each one costs you."

**What holds this together and what separates GAINS' approach from the considerable noise in the market is governance.**

Every agent operates within policy-defined boundaries. Every action is logged and auditable. Every escalation routes to a human with full context. The goal isn't AI that operates independently of human judgment. It's AI that makes human judgment faster, better-informed, and concentrates on the decisions that actually require it.

GAINS has spent years building the infrastructure that makes AI safe and valuable at scale. The result is a platform where AI isn't added on it's built in, governed, and accountable. And that, in the end, is the only foundation on which AI in supply chain can consistently create value.

## The environment has changed. The tools finally have too.

Supply chains no longer operate in a stable world—lead times swing, demand signals decay fast, and complexity outpaces what planners can manage manually. Yet most "AI for supply chain" still stops at insights (dashboards, alerts, chat) or isolated predictions (better forecasts) that don't change what gets ordered, moved, or prioritized under real constraints.

The real requirement is a decision system: that combines calibrated prediction with embedded optimization to quantify trade-offs, automates routine decisions inside policy guardrails, orchestrates across time horizons so local fixes don't create future pain, and governs every action with decision logs for accountability and continuous learning. GAINS' DEO approach delivers decision speed without surrendering control—keeping humans in the loop only where judgment is truly required, and letting the system handle the rest safely, at scale.

**To begin your journey into supply chain decision engineering and orchestration with GAINS Pragmatic AI talk to an expert at GAINSystems.com**

# Let's Move Forward Faster.

Drive Results. Engineer Your Advantage and orchestrate smarter decisions—together.

Contact GAINS today