Why We Built Autonomy

Autonomy Platform · Public Blog

CONFIDENTIAL

Company March 2026

Why We Built Autonomy

The supply chain planning industry is stuck. Here's our diagnosis, our guiding policy, and the coherent actions we're taking to fix it.

The Problem: Planning Hasn’t Changed in 20 Years

Most supply chain organizations operate the same way they did a decade ago. Planners open spreadsheets or legacy screens every Monday, review thousands of SKU-level exceptions, make judgment calls under time pressure, and hope the plan holds until next week.

When disruptions hit mid-cycle, the response is reactive: phone calls, expedited shipments, and costly overtime. The institutional knowledge that makes this work lives in the heads of a few senior planners, and when they leave, it leaves with them.

The Diagnosis

The planning function has three structural problems:

1. Periodic Cadence in a Continuous World

Supply chains don't wait for the Monday MPS run. A supplier delay on Tuesday, a demand spike on Wednesday, a quality hold on Thursday: each requires response, but the system only replans on schedule.

2. Human Bandwidth as the Bottleneck

A planner reviewing 847 exceptions per week can't distinguish signal from noise fast enough. By Thursday, she's reviewed 400. The other 447 roll into next week. Three of them were urgent. One caused a stockout.

3. Knowledge That Walks Out the Door

The rules, heuristics, and judgment calls that make planning work are undocumented and non-transferable. When a senior planner retires, 20 years of pattern recognition leaves with them.

Our Thesis

> *“ "Gartner designated Decision Intelligence 'transformational' in the 2025 AI Hype Cycle and published the inaugural Magic Quadrant for Decision Intelligence Platforms in January 2026. The shift is clear: from 'data-driven' to 'decision-centric.'" ”*

Supply chain planning is ripe for what Jordi Visser calls the “agentic inversion”: the structural shift from human labor to machine execution. Not automation (the same tasks, faster) but inversion: agents own decisions by default, humans provide governance.

We believe the transition happens through a deliberate, measured progression:

1 **Decision Support**

Human in the loop. AI recommends, human decides.

2 **Decision Augmentation**

Human on the loop. AI decides within guardrails, human inspects.

3 **Decision Automation**

Human out of the loop. Overrides during augmentation are the training signal that enables automation.

The Guiding Policy: Automate-Inform-Inspect-Override

Autonomy doesn't replace planners. It restructures what they spend their time on:

The planner becomes a manager of decisions, not a doer of tasks.

Automate	Routine decisions execute within guardrails without human involvement. Guardrails are business rules: max order value, max safety stock change, min service level floor.
Inform	Actions the system took that the planner should be aware of but doesn't need to act on.
Inspect	The agent has acted. The planner can examine any decision, reasoning, alternatives considered, calibrated confidence, at any time. The agent surfaces the ones most likely to need attention (high urgency, low confidence) at the top of the queue.

Override	Planners can always override any agent decision. The override is captured, reasoning recorded, and the system learns from the outcome.

From 847 Exceptions to 14

What happens when an enterprise planner arrives Monday morning:

612
Auto-Resolved
High likelihood, agent decided

168
Abandoned
Low urgency + low likelihood

53
Informational
Handled, flagged for awareness

14
Inspect & Override
High urgency + low likelihood

All 847 decisions were handled by the agent overnight; none waited for her approval. The 14 at the top of her inspection queue are the ones the agent flagged for review, high urgency and low calibrated confidence. She spends her morning on those: she inspects each, examines the reasoning and the alternatives considered, and overrides where her judgment differs. **Every override teaches the system for next time.**

The Self-Reinforcing Advantage

More decisions → Better AI → Less human effort → More decisions handled → ...

This loop is built on your data, your team's judgment, and your specific operating context.

The judgment layer becomes the moat.

Every agent decision generates a decision-outcome pair. The agent observes the state (inventory levels, demand signals, capacity), takes an action (order, rebalance, defer), and then measures the outcome against a balanced scorecard. Actions that improved outcomes get reinforced. Actions that didn't get dampened.

Unlike human planners who work 40 hours a week, **agents work 168**. Every hour of the night, every weekend, every bank holiday: agents are observing, deciding, and learning. They handle the repetitive and the mundane so planners can focus on the decisions that truly need human creativity.

Adoption Builds Trust Through Measurement

Week 1	~45% auto-executed decisions
Week 12	~72% auto-executed decisions
Steady State	~85% auto-executed (less than 10% overridden)

Our Approach: What We're Building

We are not building another forecasting tool or dashboard. We are building a unified decision platform for supply chain, **one** shared world model **driving** six decision domains (portfolio, demand-shaping, supply, production, transport, warehouse), all surfaced through a single Decision Stream.

- **Decisions as digital assets.** Every recurring decision is a trackable asset with defined inputs, logic, ownership, and measured outcomes.
- **Full decision lifecycle.** Model, orchestrate, monitor, and govern decisions end-to-end.
- **Research-grounded.** Every architectural decision maps to peer-reviewed research in sequential decision-making, decision science, and conformal prediction.
- **Enterprise-ready.** AWS Supply Chain data model compliance, SAP integration, RBAC, full audit trail.

11 Autonomous Agents	**24/7** Continuous Operation
<10ms Decision Latency	**35** AWS SC Entities

References

- *Tacit Knowledge Is Your Next Competitive Moat*, California Management Review (Berkeley), March 2026. The judgment layer Autonomy captures, override-by-override, is what Berkeley names the next competitive moat.
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review

Press. The economic case for treating decisions, not predictions, as the unit of value.

- Agrawal, A., Gans, J., & Goldfarb, A. (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press. Why decision systems, not point-solution AI, are the durable architecture.
- Visser, J. (2026). *The Agentic Inversion*. Substack, February 2026. The "overlap moment" framing.
- Gartner (2025-2026). *AI Hype Cycle* and inaugural *Magic Quadrant for Decision Intelligence Platforms*.