Reframing S&OP as a real-time decision engine

In markets that move daily, the binding constraint of Sales and Operations Planning (S&OP) is decision latency, not forecast accuracy.

The sharpest forecast no longer wins. In markets that move by the hour, the binding constraint on S&OP is the time between signal and committed decision.

Why decision latency is the new S&OP scorecard

- Forecast accuracy has stopped being the binding constraint on S&OP; decision latency now defines competitive advantage.
- Demand signals update hourly while supply constraints update by cycle, leaving plans consistent on paper but operationally infeasible.
- S&OP is shifting from a monthly coordination forum into a continuous decision layer powered by AI and agentic orchestration.
- Walmart, Amazon, Kroger, and River Island are already funding decision layers that compress latency from weeks to hours.

Why the monthly cadence is breaking down

S&OP exists to reconcile three truths that rarely agree on their own: what the market is actually buying, what the supply network can physically deliver, and what the business can profitably promise. Sitting at that intersection, S&OP is jointly owned by supply chain, commercial, and finance functions. It is where strategy converts into committed decisions on demand, inventory, capacity, and margin. For decades, monthly reconciliation was enough. Across retail, quick-service restaurants, logistics, distribution, and consumer packaged goods, the S&OP calendar imposed a useful discipline on planning. But the calendar has become the constraint. Volatility now lives inside the planning horizon, and a process built to align people once a month cannot keep pace with markets that move by the day. The strain on S&OP is structural, not a failure of effort. The pressure has shifted from how accurate a forecast is to how quickly, and how confidently, a decision can be made on it.

> **70% of large enterprises will adopt AI-based supply chain forecasting by 2030 (Gartner)**
>
> **91% of operations and supply chain leaders expect to significantly change their strategies (PwC, 2025)**
>
> **92% say their technology investments have not fully delivered expected results (PwC, 2025)**

The shortfall is not appetite for change; leaders are investing heavily. The gap is architectural: fragmented data, a decision cadence locked to the calendar, and weak cross-functional orchestration.

Why a sharper forecast is not the answer

It is tempting to treat S&OP as a forecasting problem and assume a better forecast will fix it. It will not. S&OP is end-to-end demand and supply planning. It has to reconcile what can be sold with what can be made, moved, and funded, all within financial guardrails. The recurring failure mode is an asymmetry of refresh rates. Demand signals update continuously through promotions, price moves, weather, local events, and channel shifts. Supply constraints update far more slowly through capacity, lead times, inbound variability, labor, and transportation. The result is a plan that is internally consistent on paper and operationally impossible in practice.

Root causes that recur consistently

- Demand, supply, and financial data sit in separate systems, so a single version of the truth always arrives late.
- Master data drifts across products, locations, and partners, quietly eroding trust in the plan.
- Planning platforms were tuned for a monthly workflow, not live decision support.
- What-if analysis remains shallow, rarely modeling a demand shift and a supply constraint together, mid-cycle.

Reframing S&OP: from forecast cycles to decision orchestration

Modernizing S&OP is less a technology upgrade than a change in what the process is for. The shift becomes clear once three stages are separated: what traditional S&OP was optimized for, what most organizations actually need today, and what leaders are already building toward.

Forecast-centric S&OP (yesterday)

When volatility was moderate and change moved slower than the planning cycle, the monthly model worked well. It optimized for consensus, reviewed variance after the period closed, and resolved exceptions in meetings rather than in workflows. The forecast was the deliverable.

Constraint-aware S&OP (today)

Volatility has crossed inside the planning horizon. Promotions, price changes, local events, weather, and channel shifts can invalidate a plan within days. Inbound variability, supplier performance, and logistics disruption move supply constraints just as fast. Most enterprises are mid-transition: reconciling demand and supply more often, running scenario reviews between cycles, and wiring exception escalation into workflows.

Decision-orchestration S&OP (tomorrow)

The destination treats S&OP as a continuous decision layer rather than a coordination forum. Planning is becoming decision-centric, with continuous sensing and recalibration replacing static plan adherence.

- Forecasts become inputs, not finished outputs.
- Supply constraints are evaluated continuously, not on a cycle.
- Decisions are simulated before execution, not corrected after they land.

The operating posture changes accordingly: from explaining variance to acting on exceptions, from defending a fixed plan to managing trade-offs dynamically, and from reconciling after the fact to rebalancing continuously. The aim is not perfect foresight. It is faster, better-founded decisions under uncertainty.

The modern S&OP decision engine: Brillio's building blocks

A real-time decision engine is not a single platform to buy. It is six connected layers operating as one value chain, with AI and agentic capability running through every layer rather than sitting on top of it. A governed data foundation holds the whole structure together.

1. **Reference and master data**
   Standard definitions for product, supplier, customer, and location form the semantic backbone. AI now watches for master-data drift, proposes harmonization rules, and catches anomalies before they propagate into the plan.
2. **Data collection**
   ERP, WMS, TMS, POS, and order systems are joined to external signals such as weather, traffic, supplier telemetry, social, and macro indices, then streamed continuously. AI-based connectors and event detection turn batch refreshes into near-real-time signal capture.
3. **Data governance**
   For AI-driven decisions, lineage, access control, compliance, and trust are non-negotiable. Governance agents monitor data quality continuously, attach confidence scores to inputs, and divert low-trust signals for review before they reach a planner.
4. **AI forecasting and decision processing**
   This is where AI becomes operational. Demand-sensing foundation models, causal forecasts that ingest external variables, and inventory-positioning models work in concert. Planner judgment is augmented, not replaced. Planners review AI-generated scenarios and approve trade-offs instead of rebuilding forecasts in spreadsheets.
5. **Workflow orchestration**
   Agentic AI earns its strongest role here. Approvals, exception triage, supplier alerts, and replanning loops are run by agents that sense an exception, simulate the options, recommend a decision, and route it to the right human. Calls that once waited for a monthly meeting now close within hours.
6. **Decision consumption**
   Outputs are built for executives to commit and operators to act, not for analysts to admire. AI-generated narratives, scenario comparisons, and recommended next actions replace static dashboards. The experience shifts from ‘read the plan’ to ‘approve the decision.’

> **AI and agents are not a layer above the building blocks; they are the connective tissue inside every one of them.**

Real-time orchestration, already in practice

None of this is hypothetical. Leading retailers are wiring real-time signals, decision intelligence, and workflow automation into their operating rhythm, and each example brings a different building block to life.

Walmart: Demand sensing wired to decision workflows

Walmart's Trend-to-Product capability uses AI and generative AI to synthesize trend signals and compress design and execution. The company reports cutting as much as 18 weeks from its traditional fashion timeline, collapsing research and design from weeks to minutes, and moving from concept to shelf in six to eight weeks. The S&OP lesson is direct: connecting demand sensing straight to decision workflows turns planning proactive, especially in short-life cycle, trend-driven categories.

Amazon: Forecasts enriched with live external variables

Amazon's foundational AI forecasting model folds in time-bound data such as weather and holidays to position inventory more precisely. The company reports a 10% gain in long-term national forecasts for deal events and a 20% gain in regional forecasts across millions of popular items. This is the forecasting layer at work: predictions continuously upgraded with external signals and tied directly to inventory positioning and execution.

Kroger: Sensing the supply side as a planning prerequisite

Kroger's supplier and carrier visibility mandates make a supply-side point plainly. Planning fails when inbound reality is unknown. The retailer now requires near-real-time location updates on inbound deliveries to its distribution centers (effective May 1, 2025) and near-real-time temperature updates (effective Nov. 1, 2025), refreshed every 15 minutes. Supply constraints must be sensed continuously, not discovered weeks later.

Nedap x River Island: inventory truth as the foundation

Item-level inventory visibility is infrastructure, not a use case. River Island reached 97% stock accuracy through item-level RFID visibility, a reminder that signal quality is the precondition for any reliable demand or supply decision. Nedap also reports a 2026 partnership with C&A to optimize omnichannel operations on the same foundation. For omnichannel retailers, weak inventory data compounds error through every downstream decision.

The macro signal: retail is funding AI-native planning

Many of the problems now attracting AI investment are not new. What changed is not just model performance. The full operating environment matured.

$113B US retail technology spend in 2026, up 6.6% year over year (Forrester)
90% of retail and CPG respondents expect to increase AI budgets in 2026 (NVIDIA)
47% are already using or assessing agentic AI in retail and CPG (NVIDIA)

Read together, these numbers point one way. Retailers are funding decision layers that cut latency, not more dashboards that describe yesterday.

Closing the decision-latency gap

The maturity debate has moved on. The question that matters now is latency: how fast can the organization sense a change, simulate the trade-offs, and commit to a decision it trusts? Every layer of the engine exists to shorten that cycle. Across enterprise engagements, four moves consistently take decision latency from weeks to hours.

1. Lead with decision velocity, not tool replacement
Set the target in time, not accuracy. Measure elapsed time from signal to commitment for a defined decision class, such as a promotion repricing, an allocation shift, or an expedite call. Choose tools to serve that metric, not the other way around.

2. Unify demand, supply, and finance on one governed foundation
Without consistent master data, governance, and shared definitions, orchestration only produces faster disagreement. The readiness gap is operating model and data, not platform.

3. Move scenarios and workflows into the daily rhythm
Take what-if simulation and exception handling out of the monthly meeting and into daily operations, where agents can route a decision to the right human in minutes.

4. Instrument decision latency as a board-level KPI
What cannot be measured cannot be improved. Track decision latency, scenario coverage, and exception-resolution speed as first-class metrics alongside forecast accuracy and service levels.

In volatile markets, forecasting explains. Orchestration decides.

> **No algorithm can outrun mismatched refresh rates**
>
> Skeptics argue that sharper forecasting models will eventually close the latency gap and that an operating-model overhaul is excessive. But no model can outrun supply truth that refreshes monthly against demand truth that refreshes hourly.

Four moves to compress S&OP decision latency

- Make decision latency the headline metric. Track time-to-commitment by decision class, not forecast accuracy alone, as the new measure of S&OP maturity.
- Unify demand, supply, and finance on one governed foundation. Orchestration without shared master data and trust only produces faster disagreement across functions.
- Move scenario simulation and exceptions into the daily rhythm. Deploy agents to triage, simulate trade-offs, and route decisions to the right human.
- Elevate latency to the board scorecard. Measure scenario coverage and exception-resolution speed as first-class KPIs; what is timed in days will compress to hours.

The next planning cycle will not be won by the best forecast, but by the shortest distance between signal and decision.

About Brillio

Brillio is The Enterprise AI Accelerator helping Fortune 1000 companies move from AI ambition to scaled impact, faster. Powered by our AI accelerator platform – Agentic Data and Application Management (ADAM), Brillio is one of the fastest-growing digital technology service providers, delivering transformation across five core workstreams: business-led transformation, customer experience transformation, AI and data engineering, digital engineering, and infrastructure engineering.

With 14 delivery locations across North America, Europe, and Asia and a team of over 6,000 customer-obsessed professionals, Brillio combines deep industry expertise, modern engineering, and accelerators to deliver measurable outcomes. Headquartered in Dallas, Texas, Brillio serves clients globally with a commitment to speed, scale, and measurable impact.

https://www.brillio.com/
Contact Us: info@brillio.com

brillio