From scheduled tasks to smart systems: The rise of Agentic AI in the supply chain

NashTech.

What's inside

The evolution of automation: This is not new!

For decades, businesses have relied on simple automated processes to keep operations ticking along. We're familiar with scheduled tasks that run like clockwork - the nightly database backup, the monthly inventory report, or the weekly sales analysis. These tools are useful but fundamentally "dumb." They do exactly what we tell them to do, nothing more and nothing less.

For example, think about a supply chain management system. It might flag delayed shipments, generate exception reports, or even send alerts when stock levels fall below predetermined thresholds. These systems are predictable and reliable, but they're also inflexible. They can't adapt to unexpected circumstances or make judgment calls when things happen outside of their original parameters.

Traditional automation executes tasks but lacks the ability to pursue goals independently. When the unexpected happens - and in the real world, it inevitably does - human intervention becomes necessary.

Enter agentic AI: A fundamental shift in approach

Agentic AI represents something quite different. **Rather than simply following instructions, these systems understand objectives and actively work to achieve them.** They possess several qualities that distinguish them from traditional automation:

Goal-orientation:
Instead of just running tasks, agentic systems pursue outcomes. They focus on the "why" behind operations, not just the "what" and "how."

Adaptability:
When circumstances change, agentic AI adjusts its approach, accordingly, finding new paths to the desired outcome.

Initiative:
Perhaps most importantly, these systems can identify when additional information or actions are needed and seek them out proactively.

Contextual understanding:
Agents grasp the bigger picture, allowing them to make smarter decisions that consider multiple factors.

To illustrate this difference, let's imagine a self-healing supply chain equipped with agentic AI. While a traditional system might flag that a shipment from Shanghai to Liverpool is delayed due to port congestion, an agentic system could go much further:

- It monitors news feeds, weather patterns, and logistics APIs in real-time.
- Upon detecting potential disruption, it calculates the likely impact on delivery schedules.
- It identifies and evaluates multiple alternative routes, modes of transport, and potential suppliers.
- It considers factors such as cost implications, contract terms, insurance cover, and carbon footprint.
- It implements changes by communicating with logistics partners through their systems.
- It adjusts internal inventory planning and notifies affected departments

All of this happens without requiring a human to initiate the process or manually evaluate options. The system understands the goal - maintaining supply chain integrity - and takes independent action to achieve it.

The complexity explosion: Why management matters

While the potential benefits of agentic AI in supply chains are enormous, so too are the challenges. The complexity of these systems can grow at a frightening pace, quickly outstripping our ability to understand or control them without proper frameworks in place.

Cascading decisions and unforeseen consequences

What begins as a seemingly straightforward decision can trigger a cascade of adjustments throughout interconnected systems. Consider this scenario:

An AI agent detects potential disruption at the Port of Rotterdam and decides to reroute shipments to Antwerp. This seems sensible enough, but consider what follows:

- The change affects arrival dates, which impacts warehouse staffing schedules.
- Different handling fees at Antwerp affect product margins across markets.
- The new route requires different documentation, triggering compliance protocols.
- Transport from Antwerp to inland destinations follows different routes, affecting delivery sequencing.
- The shift in volume affects negotiated rates with multiple logistics providers.

Before you know it, what appeared to be a simple rerouting decision has implications across operations, finance, compliance, and customer service. Without proper oversight, these cascading effects have implications that are usually the responsibility of managers in the chain.

Multi-dimensional optimisation challenges

Supply chains don't optimise for a single factor. They balance speed, cost, reliability, sustainability, and customer satisfaction simultaneously. Each factor introduces new variables and constraints.

An agentic system might identify a cheaper delivery route that increases carbon emissions, which then triggers sustainability protocols requiring carbon offsetting, which in turn affects profitability metrics and perhaps even ESG reporting requirements. These interconnected considerations create decision spaces of mind-boggling complexity, and we need to figure out who owns that space and how we contract for the liabilities.

The information explosion

As agentic systems consume more data sources - weather services, news feeds, social media, market indicators, and hundreds of other inputs - the decision complexity grows exponentially. Adding just one new data source can create entirely new decision factors that weren't previously considered.

For instance, incorporating social sentiment analysis might lead the agent to avoid routing through locations with negative public perception - a consideration that wouldn't have existed in the system's model before.
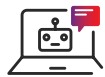
A structured approach to taming complexity

Given these challenges, how can businesses harness the power of agentic AI without being overwhelmed by its complexity? The answer lies in taking a structured, evolutionary approach that builds complexity gradually while maintaining appropriate controls.

Staged implementation: Walking before running

Rather than attempting to implement something like a fully autonomous supply chain overnight, organisations should adopt a phased approach:

Decision support phase:
The agent provides recommendations to human operators but doesn't implement changes directly

Limited autonomy phase:
The agent handles routine decisions within narrow parameters while humans manage exceptions

Supervised autonomy phase:
The agent manages most operations with humans monitoring performance

Advanced autonomy phase:
The agent operates independently with periodic human review

This gradual evolution allows both the technology and the organisation to adapt. It gives people time to understand the system's capabilities and limitations and provides opportunities to refine the agent's decision-making before critical functions are delegated.

Simulation and scenario testing

Before entrusting real operations to agentic systems, organisations should extensively test them in simulated environments. Digital twins of supply chain operations allow agents to demonstrate how they would respond to various disruptions without real-world consequences.

"War games" can test the agent's response to historical disruptions like the Suez Canal blockage or COVID-19 lockdowns. These exercises reveal how the system handles complex scenarios and identify potential weaknesses before they impact actual operations. It's not inconceivable that much of the responsible operator's work will shift to war-gaming and improving decision making while the "known" problems get solved automatically.

Robust monitoring and oversight

Traditional monitoring approaches that focus solely on outcomes are insufficient for agentic systems. Organisations need visibility into not just what decisions were made, but why they were made.

A comprehensive monitoring framework should include:

Decision rationale tracking:
Systems that record the factors considered, alternatives evaluated, and reasoning behind each decision

Balanced performance metrics: Measurements that go beyond simple efficiency to assess resilience, adaptability, and alignment with broader business objectives

Anomaly detection:
Tools that flag unusual patterns or unexpected decisions for human review

Feedback mechanisms: Channels for humans to provide input on agent decisions, which then inform future decision-making

Human-AI collaboration frameworks

Most importantly, companies need to develop clear processes for human-AI teamwork. This includes:

Protocols for humans to override agent decisions when necessary, including time for this to happen.

Communication channels for agents to explain their reasoning.

Training for staff on how to effectively work alongside AI systems.

Clear accountability frameworks that define who is responsible for what.

The path forward: Evolution not revolution

The implementation of agentic AI in supply chains isn't about replacing existing systems overnight. It's about evolutionary progress toward more intelligent, resilient operations.

By starting with well-defined use cases, implementing proper governance structures, and gradually expanding capabilities, organisations can harness the tremendous potential of these technologies while managing their inherent complexity.

The most successful implementations will be those that view **agentic AI not as a replacement for human intelligence, but as a powerful complement to it**. The goal isn't fully autonomous supply chains that operate without human involvement, but rather "augmented supply chains" where human expertise is amplified by AI capabilities.

In this model, agentic systems handle routine operations and provide decision support for complex situations, while humans focus on strategy, relationship management, and handling truly exceptional circumstances. This partnership approach maximises the strengths of both human and artificial intelligence.

Measuring success: Beyond traditional ROI

How should companies measure the return on investment from agentic AI initiatives? Traditional metrics focused solely on cost reduction or efficiency gains fail to capture the full value of these systems.

A comprehensive measurement framework should include:

- **Resilience metrics:** How quickly can the system adapt to disruptions? Beyond cost savings, successful implementations reduce "time-to-adapt"—the lag between detecting a disruption and implementing a solution—from days to minutes. Measuring recovery time from various types of disruptions provides insight into the system's true value.

- **Decision quality:** Tracking not just the speed of decisions but their quality over time. This might include measuring how often human operators accept AI recommendations versus overriding them, and whether those overrides decrease over time as the system improves.

- **Predictive accuracy:** How well does the system anticipate potential disruptions before they occur? The ability to forecast issues and take preventive action represents significant value that traditional metrics might miss.

- **Sustainability gains:** Reductions in carbon footprint, waste, or resource usage that result from more intelligent routing and resource allocation.

- **Knowledge capture:** How effectively is the organisation building institutional knowledge through the AI system? This might be measured by the system's ability to apply learnings from one disruption to similar situations in the future.

- **Opportunity cost reduction:** Time saved by management no longer dealing with routine disruptions, allowing them to focus on strategic initiatives instead.

Many of these metrics require establishing new baselines and measurement methods. For example, companies might need to document their current "time-to-adapt" for different types of disruptions before implementing agentic systems, in order to demonstrate improvement.

Conclusion: Managing the complexity is the real challenge

While the technology behind agentic AI in supply chains is sophisticated, it is increasingly accessible. **The real challenge isn't implementing these systems but managing the explosion of complexity that follows**.

By taking a structured approach that emphasises gradual evolution, clear governance, effective measurement, and human-AI collaboration, organisations can navigate this complexity successfully. The result will be supply chains that aren't just more efficient, but fundamentally more resilient and capable of adapting to disruption rather than merely responding to it.

The journey toward truly intelligent, self-healing supply chains won't be completed overnight. But for those willing to take a thoughtful, structured approach to implementation, the competitive advantages will be substantial and long-lasting.

The question isn't whether supply chains will adopt agentic AI—it's whether they'll do so thoughtfully or reactively. **The winners will be those who master the balance of autonomy and control, through a controlled and iterative implementation.**

At NashTech we have 25 years' experience of working closely with our supply chain and logistics clients to deliver sustainable technology change. **Why not talk your situation through with us, and let's find a path to the future together?**

Contact us at

NASHTECH
25 YEARS
A journey of trusted partnerships.
Delivering meaningful business transformation through technology and world-class talent.

We are experts in technology, delivering smart solutions that solve business challenges and create value. Our award-winning teams apply deep expertise and passion to deliver complex IT projects globally.

𝕏 **@NashTechHN**

in **NashTech**

Nash Tech.