

Strengthening Supply Chain Resilience Through Agentic Execution

Why AI Agents And The Knowledge
Layer Are Instrumental When
Shifting From Reactive To Trusted,
Proactive Supply Chain Success

Table of Contents

Executive Summary	03
Returning Supply Chain To A Proactive Posture	04
Realizing Agentic AI Success In Supply Chain Operations	05
AI's Knowledge Gap Blocks Supply Chain Success	06
Enterprise Knowledge Ensures Accurate, Trusted AI Outcomes	07
The Knowledge Layer Enables The Agentic Supply Chain	09
Maximizing The Supply Chain's Value: Moving From Reactive Operations To Agentic Execution	10
Ensuring Agentic Success With An Outcome-First Approach	11
An Outcome-First Supply Chain Case Study	12
The Agentic Supply Chain Is Here	13
More Resources For Supply Chain Decision-Makers	13

Executive Summary

Growing economic and geopolitical volatility dictates how enterprises must approach supply chain operations. While teams continue to juggle disruption as a cyclical event, it's now wise to accept it as the structural baseline. Supply chain chaos, unfortunately, is here to stay.

Adding more data, workers, and applications has provided little advantage, instead increasing costs, friction, and technical debt while diminishing agility. The focus must now turn to increasing supply chain resilience and certainty through AI and AI agents that leverage existing infrastructure.

However, agentic execution in supply chain use cases requires more than just data, and more than simply extending the function's existing AI-enabled foundation of predictive analytics and demand forecasting. Increasing agility at scale requires rethinking supply chain operations with modern AI agents infused with enterprise knowledge.

This Knowledge Layer, a combination of semantics, guardrails, policies, expertise, and more, delivers data intelligence, governed insights, and trusted agents for accurate, compliant, and transparent execution. Agents can then monitor forecast accuracy, supplier performance, and inventory risk in real time, using certified data, and triggering approved workflows that reduce stockouts, increase forecast precision, and protect margins.

Returning Supply Chain To A Proactive Posture

Supply chains are besieged by constant uncertainty. Originally viewed as once-in-a-generation “black swans” at the outset of the COVID-19 pandemic, stockouts, excessive lead times, staff shortages, and inflation have become familiar mainstays of the global economy. Now, those impacts have been joined by shifting trade sanctions, tariffs, and reshoring edicts.

Global conflicts such as the crisis involving Iran also arise quickly and unexpectedly, turning a regional issue into a global supply chain upheaval that blocks shipping lanes and quickly spreads the impact from oil to fertilizer and food, helium and semiconductors, and beyond.

The result is a compounding complexity of supply chain planning that nearly renders traditional linear planning models obsolete. For one component alone, U.S. tariffs, 72% of trade professionals say related volatility is “the most impactful regulatory change” they face this year, nearly doubling those saying so in the prior year. And, nearly one-quarter of supply chain organizations still rely on manual tools such as spreadsheets and email, which are incapable of providing the necessary agility, compliance, or speed. In short, reactive operations not only delay solutions, they also eat into margins and frustrate downstream operations and end customers

To truly elevate the supply chain function as a strategic partner contributing to business growth, and stop the constant swirl from fighting endless fires, agentic AI is here to automate, accelerate, and evolve supply chain operations. This paper explains how leaders can pivot from reactive to proactive postures using AI, offers an outcome-first approach that ensures durable, sustainable AI deployments, and provides use cases showing how leading organizations are already using AI to increase supply chain resilience.

Realizing Agentic AI Success In Supply Chain Operations

The enterprise potential of AI has exploded in recent years, especially with the proliferation of chatbots that can provide answers and generative AI that can create new content.

However, where generative AI and large language models (LLMs) merely summarize existing data and content, AI agents actively execute tasks such as supplier evaluations, risk monitoring, and contract analysis. Agentic systems are also capable of speeds and accuracy that far surpass human capabilities.

A few practical examples of supply chain use cases for AI agents include:

- **Inventory visibility:** AI agents capture, clean, and unite inventory data from hundreds of facilities, accounting for varied processes, formats, and systems to deliver a consolidated view. Instead of [stockpiling duplicate parts](#) in every location or blindly replenishing stockouts, teams can ask an agent and instantly know if a nearby location has the necessary parts.
- **Preventing stockouts:** AI agents monitor inventory thresholds while scanning demand forecasts and production schedules for deviations. When demand exceeds supply, agents proactively analyze supplier capabilities and trigger compliant replenishment based on cost, speed, or other factors to avoid stockouts, revenue impacts, and inventory holding costs.
- **Enhancing demand planning:** AI agents combine, analyze, and track forecast versus actual data, automatically generating comparisons, identifying biases, and surfacing potential solutions while giving analysts more time to collaborate with the business. When data quality issues arise, agents proactively curate data sources and [update metadata](#).
- **Managing supplier risks:** AI agents dynamically monitor supplier on-time, in-full (OTIF) and reliability metrics to identify and rate supplier performance. When agents detect performance issues or concerning trends, they alert human workers with recommended escalation paths or alternative sourcing solutions to avoid supply chain disruptions.

AI's Knowledge Gap Blocks Supply Chain Success

While AI's potential for supply chain use cases is massive, realizing its ROI isn't without challenges. The most significant is ensuring agents have access to enterprise knowledge. Simply [deploying general-purpose agents delivers generalized outcomes](#), not the unique outcomes required for a specific domain, business unit, or evolving enterprise.

For example, a long-time employee may have learned over time that purchase order data from Operations must be modified before being combined with similar data from Logistics, and that supplier data contains sensitive information that shouldn't be accessible to every employee. But a generic agent doesn't have this semantic, governance, and process knowledge out of the box.

It's a common challenge. As part of its data management initiatives, [Swire Coca-Cola USA found inconsistent definitions in use for critical KPIs](#) like OTIF. "[T]here were different versions of OTIF, the way OTIF was calculated, going around. The manufacturing had one or supply chain had one, OTIF is being calculated in a different way by our sales team, and so there was no consistent way of calculating it," explained Bharathi Rajan, the company's Vice President of Enterprise Data, Insights & Applications. Workers may have workarounds for these types of inconsistencies, but an agent would not have that information.

The impact of this knowledge gap is clear: Gartner estimates that organizations will [abandon 60% of AI projects](#) because they lack AI-ready data. Not just data, but AI-ready data. This means data plus contextual metadata, governance policies, decision precedents, processes, and more — known as the [Knowledge Layer](#) — are codified into [AI-ready data products](#), ensuring that agents deliver accurate, compliant, production-ready insights instead of generic guidance or misleading approximations.

Enterprise Knowledge Ensures Accurate, Trusted AI Outcomes

It's a common refrain that bears repeating: Agentic AI elevates human workers out of mundane tasks so they have more time to grapple with shifting strategic priorities, changing market dynamics, and increasing global volatility.

Think freight routing: AI agents automatically linking disparate systems and functions, making decisions that weigh options and optimize for cost, speed, or whatever is most important. Procurement, Finance, Logistics, Production, suppliers, shippers, and other stakeholders are informed, guided with best next steps, and kept in the loop for critical decisions. Products move efficiently while human workers have more time to focus on enhancing supplier relationships, advancing internal initiatives, and improving customer experiences.

However, agentic success depends on workers trusting and having confidence in agentic outputs. One critical order routed via a suboptimal shipper will quickly undermine trust in the agentic system, slowing adoption and pushing those responsible into workarounds or manual exceptions. A logistics manager might know through experience that some shippers use trucks too tall to offload at a specific customer's facility. Even if that information is in the shipper record's EQPT_RESTRICT field and the customer facility record's RECV_HT_MAX field, for example, an agent wouldn't know that out of the box, nor that the customer's data is in meters and the shipper's data is in feet.

Metadata closes that knowledge gap for agents with insights and information that makes the data AI-ready while also improving the accuracy of the outputs. Agents can then use this [Knowledge Layer](#) to validate requirements, clean and normalize data, alert responsible stakeholders when required, maintain privacy and security standards, and more.

Metadata helps agents understand data

Metadata makes enterprise knowledge portable and reusable so people, platforms, and agents can use it effectively. It explains data, semantics, policies, lineage, and more to increase accuracy and give agents the knowledge necessary to deliver trusted results.

Metadata adds context to data, such as:

- Business logic, context, and data structures for accurate interpretation.
- Data transformations to respect data flows, quality, and usage rules.
- Access controls, compliance, and governance policies based on user and data.
- Semantics to account for different meanings across domains.

In practice, agents use metadata to prevent sensitive data from being exposed, understand when to apply different governance and safety policies, and recognize usage patterns that indicate which data is more trusted or recent.

[Learn more about metadata-aware agents.](#)

The Knowledge Layer Enables The Agentic Supply Chain

While agents can easily automate generic processes, the aforementioned knowledge gaps prevent agents from delivering dependable, governed outputs consistently for domain- and entity-specific processes.

Without the necessary knowledge, many agentic pilots — some say up to 95% — fail to reach production or enterprise scale. The Knowledge Layer, however, enables confidence in agentic actions by operationalizing knowledge through the following components:

1. Data catalogs detail data assets, lineage, metadata, terms, metrics, ownership, and stewardship.
2. Data products use cataloged knowledge to deliver analytics with semantics, policies, and guardrails built in.
3. AI agents execute actions using data products plus tailored prompts and tools that combine to increase accuracy, autonomy, and confidence.

Supply chain teams use the Knowledge Layer to move reactive operations slowed by manual work and generic agents to proactive agentic execution, driven by agents grounded in enterprise knowledge. It's the key for shifting from dealing with problems after they occur to proactive, governed, agentic-enabled action.

The following table details how reactive supply chain realities leak revenue, add friction, and frustrate teams. It also shows how, with the Knowledge Layer activated, those same operations can be enhanced, accelerated, and improved through agentic execution.

Maximizing The Supply Chain's Value: Moving From Reactive Operations To Agentic Execution

Reactive operations		Agentic execution	
Reality	Outcome	Reality	Outcome
Dashboards flag stockouts at the last minute, forecast bias is discovered too late, supplier delays add up to impact customers, and manual investigations proceed slowly across systems.	Revenue leakage, excess safety stock, expedited shipping costs, firefighting culture, and margin erosion.	Agents continuously monitor forecasts and forecast accuracy against actuals, OTIF metrics, inventory thresholds, and supplier performance to detect and highlight deviations before customers and cash flows are impacted.	Proactive operational intelligence.
Unclear data lineage, multiple definitions of OTIF, fill rate, and forecast accuracy, and conflicting reports across planning, procurement, and finance teams.	Cross-functional misalignment, mistrust of analytics, slower planning cycles, and executive reporting conflicts.	Agents operate on certified data products with standardized definitions, clear ownership, and full lineage transparency.	Execution aligned across the enterprise.
Forecasting models lack awareness of supplier contracts, service-level policies, escalation workflows, and access entitlements.	Risky AI recommendations, policy violations, compliance exposure, and a loss of trust in AI outcomes.	Agents embed enterprise definitions, policies, and entitlements into every AI-driven action.	Enterprise-grade, policy-aware AI.
Planners reconcile ERP, WMS, TMS, and spreadsheets manually, cross-functional meetings waste time determining cause of variances, and diagnosis is only possible with internal SMEs.	Delayed corrective action, recurring disruptions, and a high analyst burden that exacerbates knowledge loss when employees leave.	Agents perform guided root cause analysis across certified supply chain data products, preserving insights and enabling continuity.	Continuous improvements to data happen autonomously.
AI recommendations lack audit trails, have unclear accountability, and expose potentially sensitive supplier or contract data.	CISO resistance, data governance pushback, and regulatory risk combine to stall AI initiatives before they can scale.	Every agentic action is policy-aware, entitlement-preserving, auditable, and governed with human oversight.	Controlled, scalable AI deployment.
Insights are trapped in slide decks, emails, and analysts' memory, forcing teams to treat each disruption as a new event.	Repeated mistakes, slower recovery from recurring issues, and enterprise knowledge erosion.	Agents learn from every action and interaction, strengthening the Knowledge Layer and improving subsequent monitoring and investigation.	Enterprise knowledge continually compounds.

Ensuring Agentic Success With An Outcome-First Approach

The value agents can bring to supply chain operations is well understood. A recent survey found that 70% of supply chain organizations are using AI for predictive analytics, real-time decision support, and supplier monitoring. However, just 23% have a formal AI strategy, and less than half “strongly agree” that they’ll roll out agentic AI before 2028.

As Gartner revealed, many AI initiatives delays and abandonments are attributed to a lack of access to AI-ready data and knowledge. On the contrary, AI success typically follows a clear pattern: Begin with the desired business outcome, not the data or AI model. The following is a simplified example.

Component	Reasoning
<p>Goal Reduce stockouts by 20%.</p>	<p>Beginning with the business goal ensures related current processes are accurately mapped. Decision points, required knowledge and data, how exceptions are handled, and other key process components can then be captured.</p>
<p>Agents To achieve the goal, tailored agents continuously monitor stock, predict demand, and recommend governed actions.</p>	<p>The business goal creates the foundation for defining what agents will take on, where human oversight is required, and where agents can be deployed.</p>
<p>Data products To act, the agent requires access to trusted inventory, supplier, and contract metrics.</p>	<p>Next, data products can be crafted to deliver data in accordance with policies, governance, and other enterprise knowledge that guides data acquisition and usage.</p>
<p>Data catalog and data To ensure control and accuracy, certified, governed data forms the foundation of the outcome-driven solution.</p>	<p>Finally, the necessary data can be identified and operationalized into data products that agents use to reason, act, and make decisions.</p>

With an outcome-first approach, agents can be built and deployed confidently because they’re grounded in enterprise knowledge. As they work, agents can also capture feedback and decision context that can be fed back into data products and the data catalog for continuous improvement.

An Outcome-First Supply Chain Case Study

A leading transportation equipment manufacturer in North America provides a compelling example of the power of agentic AI in supply chain operations. Knowing that higher customer satisfaction leads to higher repeat sales, the company sought to speed up delivery issue resolution times. It was a highly manual workflow that relied on data siloed across DB2, Alteryx, and mainframe systems, making product lifecycle tracking and delivery issue identification slow, frustrating tasks.

Using an outcome-driven strategy, the company focused on the desired results, and then developed AI-ready data products to provide the solution. Using its Knowledge Layer to inform AI agents, it unified order, build, delivery, and service data and delivered it to stakeholders in a cohesive strategic view.

Goal: Improve customer satisfaction to increase repeat sales

Agent: Provide product and delivery issue data in a concise report

Product: Serve up critical data so agents and workers can use it confidently

Data: Order, build, delivery, and service data

Today, teams also [chat with those data products using natural language](#) to gather real-time supply chain intelligence, directly improving delivery performance and strengthening the customer relationships that drive repeat sales.

This transition shifted supply chain operations away from a dependency on organizational memory and brittle data integrations toward a state of explainable, repeatable decision-making that scales while delivering desired business outcomes.

The Agentic Supply Chain Is Here

Most supply chain teams rely on dashboards that highlight problems after they've already impacted revenue. That's untenable in today's hyper-volatile global supply chain environment.

AI agents are here to automate operations, improve resilience and speed, and elevate human workers to more strategic, cognitive roles. Organizations simply need an AI-ready data foundation to deliver governed data, enable certified data products, and unleash the enterprise Knowledge Layer to autonomous AI agents.

[Alation provides the Knowledge Layer](#) to deploy governed AI agents that monitor forecast accuracy, supplier performance, and inventory risk in real time. These agents — built, deployed, and monitored by you — analyze certified data and trigger approved workflows, reducing stockouts, improving forecast precision, and protecting margins, all within enterprise governance and oversight.

[Book a demo today](#) to learn more about Alation's agentic solutions for supply chain operations.

More Resources For Supply Chain Decision-Makers

- [How To Use Metadata To Navigate Supply Chain Volatility](#)
- [How Data Management Can Protect Your Supply Chain From Tariffs And Disruptions](#)
- [How Daimler Trucks Uses AI Agents & Metadata To Transform Manufacturing](#)
- [How Lipton Uses Metadata & Agentic AI To Transform Operations](#)
- [How Brambles Uses CDEs To Govern Global Supply-Chain Data](#)
- [How To Automate Supply Chain Management \(Video\)](#)
- [Alation Test Flight: Building A Supply Chain Data Product](#)