RESOURCES

The Context Puzzle



The architectural approach to creating multienterprise context for an actionable agentic supply chain

1. Agentic AI Exposes the Supply Chain Context Gap

Supply chains across industries have spent years chasing agility and visibility they still do not have. The pressures differ by sector, but the dependency on external companies is shared: pharmaceuticals rely on contract manufacturers and distributors; automotive and aerospace on multi-tier suppliers; food and beverage on ingredient suppliers, logistics partners, and channels. Deloitte's 2026 outlook for 280 life sciences executives illustrates the urgency in one highly interdependent sector: 44% name pricing and access, 39% geopolitical tension, 38% inflation and supply chain risk, and 37% competition from generics and biosimilars.[1]

In that life sciences survey, 80% say future competitiveness depends on how effectively they use AI, and 78% expect it to drive major change. But reality is that only 22% say they have successfully scaled AI, and 9% report significant returns.[1] These are sector-specific results, but the architectural question extends to any multienterprise supply chain.

The shift toward agentic AI as a lever for growth, cost reduction, and margin

© TraceLink Inc. 2009 - 2025 All Rights Reserved

improvement is therefore a natural one. Making supply chain a priority for these initiatives goes beyond the argument that much of the value-add is sitting outside the four walls: the work is high-volume and still largely manual. It is measurable per transaction, which means value can be shown against a baseline. And a material share of the context required to perform end-to-end supply chain work sits throughout the multienterprise network of partners like suppliers, logistics providers, and distributors, as well as across multiple enterprise systems of record.

The dependency is particularly visible in industries with outsourced manufacturing and distribution. In life sciences, for example, a company can depend on external partners for manufacturing status, batch release, channel inventory and delivery commitments. In automotive, food and beverage, or aerospace, comparable gaps arise around supplier capacity, material availability, quality, and shipment milestones.

2. The Agentic AI Opportunity Faces Structural Barriers

Within four years, 74% of leaders expect nearly half of their business processes rebuilt around AI agents.[3] The outcomes behind that expectation are documented. McKinsey reports a pharmaceutical company that put agents on invoice-to-contract compliance and cut leakage by 4%, and an aircraft OEM that cut active inventory by 30% using agents on order execution.[4]

The barriers are documented too, and they are not about models. Among organizations already piloting agentic AI, 72% lack unified, accessible data and 67% say integration is too costly and complex.[3] Gartner predicts over 40% of agentic AI projects will be canceled by the end of 2027.[5] McKinsey finds "eight in ten companies cite data limitations as a roadblock to scaling agentic AI."[6] The evidence consistently points to data readiness and linking, not to model quality.

© TraceLink Inc. 2009 - 2025 All Rights Reserved

3. Internal Data Readiness Is a Must, but Still Leaves a Multienterprise Gap

The prescription that follows is sound: treat data as a managed asset, with quality management, metadata, lineage, and governance.[7] When organizations have multiple systems of record for the same objects, building this layer of context is one requirement for action-performing agents that go beyond analytics. For some internal operations where the data is yours, such as HR and IT, the process and the data estate can be closely aligned.

Supply chain is where they are not. A purchase order is agreed between two companies. A production commitment involves a supplier or contract manufacturer. Channel inventory sits on a distributor's system. Much of any end-to-end process is performed by parties with their own schemas and their own systems of record. That is the Context Ceiling: the limit on what an agent can do when the trusted operational context available to it stops at the enterprise boundary while the business process continues beyond it.

An agent reasons over the union of what it can contextualize, so if half the process is invisible, that union is capped however good the visible half becomes. Internal context is one of two required inputs; the internal programme is not wasted work. But the remaining gap does not close by doing more of the same, because the missing information is not yours to originate or control at its source. You cannot impose your governance model on a partner's ERP. What you can do is govern how that information is identified, exchanged, interpreted, and used once it participates in a shared process.

4. Traditional Partner Integration Limits Digital Reach

Deloitte is explicit that master data management works by "ensuring that data is FAIR, findable, accessible, interoperable, and reusable, ***across organisational***

© TraceLink Inc. 2009 - 2025 All Rights Reserved

and ecosystem boundaries."[8] Take channel inventory. You know what you shipped to the distributor. You do not know what is on their shelf, at which site, or how fast it is moving. The data is not poor. It does not exist on your side.

Closing that gap partner by partner was the default, and it is what the 67% who call integration too costly are describing. Each partner is a project, and each new business transaction type is a new project against every partner already linked. It also fails a test budgets cannot fix: the work sits on the partner's side, on their IT capacity and release calendar, so the schedule is set by the least available party. Doubling your integration budget does not double a smaller supplier's ability to respond.

There is a second constraint, and it sits on your side of the boundary. The team that would run those projects could be the same team already committed to the data lake migration, the ERP upgrade, and the planning platform rollout. Traditional partner integration does not queue behind budget; it queues behind available resources. Figure 1 illustrates how constraint compounds as more trading partners are linked. The cumulative impact as partner count increases is illustrated in Figure 1.

© TraceLink Inc. 2009 - 2025 All Rights Reserved

FIGURE 1

Cumulative time to link trading partners

Elapsed days to reach N linked partners. The gap compounds with every partner added.



The gap is not a discount. It widens with partner count.

About half the elapsed time at five partners, about two thirds at thirty-five.

Modeled by the publisher of this paper from its own customer programme data. Interested-party evidence; method and sample not independently verified.

Figure 1. Cumulative elapsed days to reach N linked trading partners. Modeled by the publisher; interested-party evidence.

At five partners the two approaches are within a few months of each other. By thirty-five, the difference is most of a working year, because on a shared network where each party integrates once to a canonical model, a partner who is already present is a configuration step rather than a project.

That is also what changes the demand on your own people. Integrating once means one integration to build rather than one per partner. And a maintenance load that does not grow with partner count, because a partner’s schema change is absorbed by the network rather than landing back in your backlog, and reach stops depending on capacity.

Two conditions decide whether reach is widely achievable, and neither is a property of your middleware. The partner has to be able to participate. A manufacturer’s network is not uniformly digital: some suppliers run modern ERP, some distributors work in spreadsheets and email. What matters is whether the same business transaction can arrive as X12, EDIFACT, an IDoc, EPCIS, XML, CSV,

© TraceLink Inc. 2009 - 2025 All Rights Reserved

an API call or a form somebody fills in, and land as the same business object either way. And the partner has to have a reason to say yes. A bilateral link asks a supplier to absorb cost for one customer; a shared link they can reuse with their other customers changes that calculation.

Reach also has to extend past procure-to-pay to the operational transactions and processes that expose emerging risk. Orders, acknowledgements and invoices tell you what was agreed and what moved. Forecasts, forecast commitments, inventory balances and production or shipment milestones are what let an agent see a problem coming. That is the multi-system, multienterprise context needed.

5. Digital Reach Must Mature into Trusted Operational Context

Reaching a partner's data is not the same as understanding it. The evidence sits inside every procurement function already: McKinsey estimates they use less than 20% of the data available to them, data they have, unused because nobody established what it relates to or how it bears on the decision.[4] A partner file arriving with no metadata-driven context will not do better.

Return to channel inventory. The weekly file lands on time, and an agent still cannot act on it. Which of your products does the SKU refer to? Which batch, expiring when? How much is already committed? Across the boundary the contextual work has to happen in the exchange itself: a canonical model that every partner format translates into, plus metadata recording what each business object is, where it came from and what it relates to. Validation belongs there too, against schema rules, business rules and partner-specific constraints, because an architecture that validates after landing has already let an agent act on a bad number.

Data alone is not the requirement. The requirement is enough trusted operational

© TraceLink Inc. 2009 - 2025 All Rights Reserved

context, internal and external, to reason and perform work correctly, and that is a ladder with five rungs:

Stage	The question a CIO should be able to answer
Reach	Can I obtain the signal from the trading partner or system at all?
Meaning	Do I know what it represents and how it relates to other business objects?
Context	Can I place it within the process, the policies, the relationships, the history and the current operating state?
Reasoning	Can a person or a governed agent weigh that context against objectives, rules and available actions?
Work	Can the resulting action be performed through a governed and auditable process?

Most programmes stop at the first rung and assume the rest follows. The third rung is where most of the difficulty sits, because process state, what is approved, what is in exception, who acted and what they decided, mostly lives in email today, which means it does not exist as data at all. Once the process state becomes data, every change in it is an event that an agent can act on.

Deloitte states the consequence plainly: without trustworthy master data, AI "is prone to risk of hallucination and incorrect recommendations."[8] Payload transport is mature. Preserving meaning, provenance, validation and shared process state across a boundary, at network scale and to a standard that survives an audit, is a different problem.

6. Enterprise Architecture Must Support Multienterprise Work

Most companies are not entering the agentic era with a blank estate, and the right posture is additive rather than replacement. To see where the gap is, start from what each part was built for.

Systems **of record** maintain authoritative enterprise state: orders, inventory

© TraceLink Inc. 2009 - 2025 All Rights Reserved

accounting, financial posting, master data. The financial controls that depend on it are a reason to leave it exactly where it is. A coordination layer does not replace that role. It links business transactions and operational signals across those systems and trading partners while preserving their business meaning and relationships. Together with turning changes into business events, this creates the shared operational context required for reasoning and work across the multienterprise supply network.

Architecture	Its primary role	Additional capability required for multienterprise agentic work
Enterprise systems of record SAP, Oracle, Dynamics	Maintaining authoritative enterprise transactions and state	Governed context spanning multiple independent companies, systems and shared processes
Supporting Enterprise Systems TMS, WMS, Planning	Operation-specific business rules and execution	Governed context and business objects across systems and shared processes
Cloud data platforms Snowflake, Databricks	Data engineering, analytics and AI or ML compute	Persistent trading-partner reach, supply chain business semantics, transaction relationships, provenance and shared process context
iPaaS and middleware MuleSoft, Workato	Moving and transforming information among applications	Persistent business semantics, network identity, shared transaction state, multienterprise processes and agent governance

The same principle applies twice more.

Deterministic automation does not become obsolete. Workflow automation, RPA, rules engines, scripts and integration flows remain the right mechanism when inputs are known, rules are stable and the action can be predefined. Agentic work can move beyond that: interpreting changing operational context, reasoning across multiple business objects and relationships, weighing governed alternatives and determining the appropriate next action. **Nor does every agent need to**

© TraceLink Inc. 2009 - 2025 All Rights Reserved

come from one place. Enterprises will increasingly run heterogeneous agent estates, with agents inside enterprise applications, function-specific platforms and multienterprise operating environments. The architectural question is whether each agent has governed access to the context, permissions, business objects and actions required for the work it is expected to perform.

Assembling the coordination layer yourself can be a legitimate answer. But be precise about what it commits you to: the partner network, the canonical transaction model, the validation and exception logic, the collaboration layer, the shared partner view and governed read/write access to every system holding the same objects, plus the event model that triggers work. The question is not whether your architecture needs them, but whether building and running them is your business.

The **Agentic Supply Chain Operating Model (ASCOM)**, **Agentic Supply Chain Operating System (ASCOS)**, and **Integrated Agentic Business Processes (ABPs)** establish the framework for governed work among humans and agents. Their value at scale depends on trusted operational context extending beyond enterprise boundaries to the trading partners and systems involved in the work.

7. CIOs Must Address the Multienterprise and Multi-system Context Gap

Establish dependency on your ecosystem. Of the process steps in your two highest-volume end-to-end processes, what fraction is performed by another legal entity? Which external signals would a governed agent need to interpret those processes reliably? How many systems of record hold the objects those steps touch?

Take a high-value, highly feasible process and baseline it before you start. Map candidate processes by their business value and their

© TraceLink Inc. 2009 - 2025 All Rights Reserved

readiness for agentic work, including the digital maturity of the process, the availability of partner information, and how easily the missing context could be established. This gives the AI roadmap a measurable starting point.

Then put five questions to any architecture:

1. **Trace.** Can it follow a business transaction across multiple systems of record and multiple companies and explain its state and provenance?
2. **Reach.** Can it link the necessary trading partners without another bilateral project for every relationship and transaction type, including partners with no EDI capability at all?
3. **Meaning.** Does it preserve the business meaning and relationships of the transaction, or does it primarily move payloads?
4. **Governance. Are agents, humans, systems, access, decisions, and actions governed through identity, permissions, policy, and auditability?**
5. **Coexistence and response.** Can it work with the systems of record, automation and agents you already run, while turning trusted multienterprise context from events into governed reasoning and work?

8. Scaling Agentic AI Requires Trusted Context Beyond the Enterprise

The gap between AI ambition and measurable return is already visible in life sciences: 80% of surveyed leaders say future competitiveness depends on effective AI, while 9% report significant returns.[1] Other industries face different pressures, but the same architectural question: can the environment extend trusted operational context beyond the four walls, preserve systems of record and AI investments already in place, and give humans and governed agents an environment to perform work?

© TraceLink Inc. 2009 - 2025 All Rights Reserved

The Context Puzzle is the boundary of the trusted operational context available to an agent. For CIOs, the question is whether their architecture can extend that context between systems to the partners and processes on which outcomes depend—and turn it into governed, measurable work.

Figure 1 is modeled by the publisher from its own customer programme data and is marked as interested-party evidence. Results and elapsed-time assumptions have not been independently verified.

Sources

1. Lyons, P., Konersmann, T., Jacobson, S., Kleyn, N., Rekhraj, K., Gosalia, D., "2026 life sciences executive outlook," Deloitte, 9 December 2025. Survey of 280 C-suite executives (180 biopharma, 100 medtech), US / Europe / Asia, fielded August to September 2025, with 9 in-depth interviews. **https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html**. *Consulting survey with disclosed method (credible secondary). At approximately 12 months since fielding, the oldest survey used here.*
2. "2025 Annual Drug Shortages Report," US Pharmacopeia, published 9 June 2026, data as of 31 December 2025. USP's own report pages returned access errors and were not read directly. The 44% single-country key-starting-material figure is quoted from USP's press release; the duration figures are from the report body as reported by RAPS and by CIDRAP at the University of Minnesota. The duration measures the age of shortages still active at year-end, not the length of a typical shortage episode. **https://www.cidrap.umn.edu/resilient-drug-supply/us-drug-shortages-have-declined-now-drag-average-5-years-report-shows**. *Standards-body report, read through secondary reporting rather than directly.*
3. "AI agents are only the beginning: The path to agentic transformation," Deloitte, 12 August 2026. Survey of 501 senior manager to C-suite

© TraceLink Inc. 2009 - 2025 All Rights Reserved

respondents in the United States, fielded April to June 2026. Every participating organization was required to be at least piloting agentic AI, so the sample is committed adopters rather than general enterprise. **https://www.deloitte.com/us/en/about/press-room/deloitte-survey-examines-ai-readiness-agentic-ai-success.html**. *Consulting survey with disclosed method (credible secondary). Figures independently re-verified.*

4. Mittal, A., Belotserkovskiy, R., Liakopoulou, T., "Redefining procurement performance in the era of agentic AI," McKinsey & Company, 5 February 2026. **https://www.mckinsey.com/capabilities/operations/our-insights/redefining-procurement-performance-in-the-era-of-agentic-ai**. *Consulting research (credible secondary). Named client cases with no sample or method disclosed; the under-20% data utilisation figure is stated in the source as a McKinsey estimate, not a measurement. Figures independently re-verified.*
5. "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," Gartner, 25 June 2025. **https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027**. *Analyst prediction (credible secondary), a forecast rather than a measurement, published as a press release rather than research.*
6. Tavakoli, A., Goodman, B., Rowshankish, K., Kumar, A., Soller, H., and others, "Building the foundations for agentic AI at scale," McKinsey & Company, 2 April 2026. **https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/building-the-foundations-for-agentic-ai-at-scale**. *Consulting research (credible secondary). Quoted verbatim and independently re-verified.*
7. Tavakoli, A., Goodman, B., Rowshankish, K., Reddin, S., "AI data readiness: the key to scaling impact," McKinsey & Company, 23 June 2026. **https://www.mckinsey.com/capabilities/mckinsey-technology/our-insi**

© TraceLink Inc. 2009 - 2025 All Rights Reserved

ghts/ai-data-readiness-the-key-to-scaling-impact. *Consulting research (credible secondary). Independently re-verified.*

8. Sinha, N., Freund, A., Andrensek, J., Roy, A., “Why Master Data Management is ‘THE’ Prerequisite for Enterprise AI Enablement,” Deloitte Switzerland, 1 October 2025. **https://www.deloitte.com/ch/en/Industries/life-sciences-health-care/perspectives/master-data-management-for-enterprise-ai.html**. *Consulting perspective (credible secondary), qualitative, contains no statistics. Quotes independently re-verified verbatim and retain the source’s British spelling.*
9. Figure 1: TraceLink internal customer programme data, modeled, 2026. *Vendor-measured and vendor-modeled. Method, sample and the definition of elapsed time are not published and have not been independently verified. Interested party.*

.tl-zoomable { max-width: 100%; height: auto; cursor: zoom-in; transition: transform 0.2s ease, box-shadow 0.2s ease; } .tl-zoomable:hover, .tl-zoomable:focus-visible { transform: scale(1.03); box-shadow: 0 6px 24px rgba(0, 0, 0, 0.25); } .tl-zoom-overlay { position: fixed; inset: 0; z-index: 99999; display: flex; align-items: center; justify-content: center; padding: 24px; background: rgba(0, 0, 0, 0.85); cursor: zoom-out; } .tl-zoom-overlay img { max-width: 100%; max-height: 100%; width: auto; height: auto; object-fit: contain; box-shadow: 0 10px 40px rgba(0, 0, 0, 0.5); } @media (prefers-reduced-motion: reduce) { .tl-zoomable { transition: none; } }
BlogAgentic Orchestration

Related Content

© TraceLink Inc. 2009 - 2025 All Rights Reserved



The Next Competitive Advantage in Supply Chain Starts With a New Operating Model

Download the white paper to see how the Agentic Supply Chain Operating Model transforms multienterprise supply network performance and productivity through governed human-agent work.

View More

© TraceLink Inc. 2009 - 2025 All Rights Reserved



Scaling Enterprise AI Starts With the Right Operating System

Download the white paper to see how the Agentic Supply Chain Operating System, built on the OPUS Platform, implements and scales the Agentic Supply Chain Operating Model through governed OPUS Agents.

View More

© TraceLink Inc. 2009 - 2025 All Rights Reserved



Agentic Orchestration: A Powerful New Operating Model for the Life Sciences Supply Chain

Agentic orchestration represents a fundamental shift in how life sciences supply chains operate—moving from fragmented processes and reactive decision-making toward more coordinated, intelligent execution across the end-to-end ecosystem.

View More

© TraceLink Inc. 2009 - 2025 All Rights Reserved