Proceedings of the 11th North American International Conference on Industrial Engineering and Operations Management, Milwaukee, Wisconsin, USA, June 9-11, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA11.20260238
Published: June 9, 2026

Why Agentic AI Stalls in Supply Chains: Insights from 32 Industry Leaders and a Framework for Foundational Enablement

Prabhat Rao Pinnaka
Supply Chain & AI Practitioner
USA
prabhatraopinnaka@gmail.com

Sukanya Bollineni
USA
Sukanyabollineni17@gmail.com

Abstract

Enterprise investment in supply chain artificial intelligence has grown substantially, yet autonomous, multi-step agentic AI systems remain at pilot stage in the vast majority of organizations. This paper investigates the gap between investment and adoption through a structured qualitative study: protocol-guided conversations with 32 purposively sampled supply chain executives, directors, and senior managers across retail, manufacturing, logistics, and consumer goods sectors, conducted throughout late 2025 and early 2026 using a consistent four-domain conversation protocol and analyzed using reflexive thematic analysis following Braun and Clarke (2022). Three foundational barriers emerged with high cross-industry consistency. First, the absence of shared domain ontologies prevents agents from reasoning coherently about supply chain entities defined differently across source systems. Second, the lack of persistent, relational knowledge structures deprives agents of the contextual intelligence that experienced human planners carry as tacit knowledge. Third, the absence of live decision context across agent cycles produces stateless systems incapable of multi-step coordination. These barriers are sequential dependencies: each layer of deficiency compounds the one before it, and organizations that skip foundational infrastructure investments consistently encounter the highest remediation costs. In response, this paper introduces SCALE (Supply Chain Agentic Layer Enablement), a three-layer infrastructure framework addressing each barrier through an Ontology Layer, a Knowledge Graph Layer, and a Context Graph Layer. SCALE reorients the supply chain AI conversation from model capability to system readiness — the precondition that existing literature has systematically underexamined.

Keywords
Agentic AI adoption, Supply Chain Management, Knowledge Graphs, Ontology, Context graphs, Qualitative research, AI infrastructure, Thematic analysis

1. Introduction
The supply chain AI investment cycle of the early 2020s produced a persistent paradox. Enterprise spending on AI and machine learning for supply chain applications has grown sharply, yet the gap between investment and autonomous deployment remains wide. Gartner (2025a) reports that just 23% of supply chain organizations have a formal AI strategy in place, and that the dominant deployment pattern is project-by-project short-term pilots rather than transformational investment. A separate Gartner (2025b) prediction warns that 60% of supply chain digital adoption efforts will fail to deliver promised value by 2028, primarily due to insufficient foundational infrastructure investment. Agentic AI — systems that take multi-step, goal-directed actions rather than generating passive

recommendations — describes both the most commercially anticipated and the most persistently under-adopted category of supply chain AI. Understanding why is the central question of this paper.

The existing literature offers two main explanations. Technology Acceptance Model (TAM) research attributes low adoption to individual-level barriers: insufficient trust, low perceived usefulness, and inadequate ease of use (Davis 1989; Ibrahim et al. 2025). Practitioner intelligence emphasizes organizational factors: change management resistance, AI literacy gaps, and data quality problems (Gartner 2025a; Hangl et al. 2023; Shahzadi et al. 2024). Both explanations locate the problem in human or organizational factors surrounding the AI system. Neither seriously examines the possibility that AI systems themselves lack the foundational infrastructure to function reliably in real enterprise supply chain environments and that practitioner resistance is a rational response to systems that do not work, rather than an irrational barrier to systems that do.

This paper investigates that architectural explanation through practitioner intelligence gathered through professional networks and industry forums. The author conducted 32 protocol-guided conversations with purposively sampled supply chain executives, directors, and senior managers throughout late 2025 and early 2026, systematically probing their organizations' AI adoption trajectories, the specific points at which deployments stalled or failed, and their own diagnosis of root causes. Thematic analysis produced three themes of high cross-industry consistency — converging on foundational infrastructure deficits rather than individual or organizational barriers.

The paper makes three contributions. First, it provides qualitative practitioner evidence focused on agentic AI adoption barriers in supply chain management. Second, it introduces a taxonomy of foundational barriers organized around ontology, knowledge infrastructure, and context — reorienting the adoption conversation from individual acceptance to system readiness. Third, it proposes SCALE (Supply Chain Agentic Layer Enablement), a prescriptive three-layer infrastructure framework addressing each barrier category with a dedicated infrastructure component and a specified organizational investment.

A note on positionality and scope. This study was conducted by a practitioner-researcher with direct accountability for AI and analytics product deployments across a Fortune 50 retail supply chain network, and prior consulting experience with AI implementations across consumer goods, chemicals, and manufacturing sectors. SCALE draws on both the practitioner insights and deployment experience; it is not validated through a controlled field study, and the paper does not claim that it is.

2. Background and Related Literature

Three bodies of literature provide the conceptual context for this study: the state of agentic AI in supply chains, the academic treatment of AI adoption barriers, and the role of ontologies and knowledge graphs in enabling AI reasoning.

2.1 Agentic AI in Supply Chains: Capability Without Deployment

Agentic AI systems are distinguished from conventional ML by their capacity for multi-step, goal-directed action without continuous human instruction (Xu et al. 2024). In supply chain contexts this means systems that respond to demand signals, exceptions, or disruptions by taking coordinated action - placing orders, rerouting shipments, reallocating inventory, triggering escalation workflows - rather than surfacing recommendations for human execution. Li et al. (2023) showed that LLMs can interpret and act on supply chain optimization outputs. Jannelli et al. (2025) demonstrated that LLM-based multi-agent consensus-seeking reduces bullwhip amplification relative to human planners. Jackson et al. (2024) mapped generative AI capabilities across thirteen supply chain decision domains, identifying autonomous planning, procurement, and exception resolution as high-potential applications. Gartner (2025c) predicts that by 2030, 50% of cross-functional supply chain management solutions will use intelligent agents to autonomously execute decisions. Walmart, Amazon, and DHL have each deployed specialized AI agents for demand sensing, fulfillment optimization, and logistics monitoring respectively, demonstrating that agentic deployment at scale is technically achievable (SAP 2025; World Economic Forum 2025).

Enterprise deployment has not followed at the pace investment would suggest. Xu et al. (2021) documented the stagnation of traditional multi-agent systems despite two decades of development; that gap persists for newer agentic systems. Gartner's 2025 Hype Cycle for Supply Chain Strategy places generative AI in the trough of disillusionment, with scaling and legacy integration challenges tempering earlier excitement (Gartner 2025d). EY (2025) reports that while many organizations have launched generative AI supply chain initiatives, 62% reassessed those programs within 12 months and only 7% completed deployment with data quality and data access cited as the top barriers by 38% and

33% of respondents respectively. AlMahri et al. (2025a) represent a partial exception with a deployed agentic disruption monitoring system, though the authors characterize their deployment as early-stage.

2.2 Adoption Barriers: The Attitudinal Literature and Its Limits

TAM-based research has produced a well-developed account of individual-level adoption barriers. Davis (1989) established perceived usefulness (PU) and perceived ease of use (PEOU) as primary adoption determinants. Ibrahim et al. (2025) validated an extended TAM in the context of AI adoption specifically, confirming that PU and PEOU remain robust predictors of behavioral intention even when personality constructs are added. Queiroz et al. (2025) applied TAM directly to supply chain AI adoption, finding that PU and PEOU are significantly associated with supply chain learning outcomes across 206 practitioners. Hangl et al. (2023) conducted a tertiary study synthesizing AI adoption barriers in SCM across 44 literature reviews, identifying change management, technical limitations, and human acceptance as the three dominant barrier categories.

Its limitation, for the purposes of this paper, is the implicit assumption that the systems being evaluated are capable of functioning as intended. Kosasih et al. (2023) argued that explainability failures in supply chain AI are primarily structural: without domain knowledge encoding, AI outputs cannot be connected to the relational context practitioners use to evaluate decisions. Brintrup et al. (2025) showed that causal ML outperforms black-box models in supply chain risk contexts not because of algorithmic superiority but because it encodes the causal structure practitioners already understand. This paper tests the proposition that system architecture — not practitioner attitude — may be the binding adoption constraint, using TAM's three core constructs (PU, PEOU, and Trust) as an analytical lens. Rather than measuring TAM constructs through a survey instrument, this paper argues that the three identified infrastructure deficits are the architectural preconditions of PU, PEOU, and Trust respectively and that TAM-based adoption interventions will fail when those preconditions are absent.

2.3 Ontologies, Knowledge Graphs, and the Infrastructure Gap

Ontologies provide formal, shared specifications of domain concepts and their relationships — the semantic substrate on which multi-agent reasoning depends (Gruber 1993). Supply chain ontology development has been active but fragmented. The SCOR model provides a widely adopted process taxonomy but lacks the formal logical axioms required for machine reasoning (Noy and McGuinness 2001). Academic ontologies have been proposed for subdomains including logistics, procurement, and sustainability, but have not converged into a shared upper ontology capable of supporting cross-functional agent coordination (Kluza et al. 2022; Bao et al. 2023).

Knowledge graphs extend ontologies with instance-level data: actual suppliers, products, facilities, contracts, and transactions encoded as typed, queryable relationships (Ji et al. 2021). Kosasih and Brintrup (2022) demonstrated their value for supply chain risk reasoning; Xu et al. (2025) showed that LLM agents operating over supply chain knowledge graphs substantially outperform agents without relational context. Kosasih and Brintrup (2024) further showed that neurosymbolic approaches combining graph reasoning with neural prediction achieve trustworthy link prediction with interpretable reasoning paths. AlMahri et al. (2025b) extended this work to relationship prediction using generative AI, demonstrating that knowledge graphs with ontological foundations enable supply chain visibility information unavailable through conventional data architectures. Despite these results, enterprise-scale supply chain knowledge graphs remain rare — the infrastructure investment required is substantial and is not treated as standard supply chain IT architecture.

The concept of context graphs extends knowledge graphs further by adding operational currency: the live record of agent decisions, reasoning traces, inter-agent communication, and human override history accumulated during operation (Atlan 2024; Foundation Capital 2023). Zero100 (2025), a supply chain technology intelligence firm, identifies conflating knowledge graphs and ontologies as a critical failure mode in AI implementation. This paper provides the empirical grounding through practitioner insights and the architectural specification through SCALE - that the context graph concept currently lacks in academic literature.

3. Research Method

This study employs a structured qualitative research design. The research question — what foundational barriers prevent agentic AI from functioning reliably in enterprise supply chain environments? — is explanatory and theory-generating rather than confirmatory. Qualitative inquiry is the epistemologically appropriate choice: the objective is to generate conceptual categories grounded in practitioner experience, not to measure the statistical distribution of known constructs (Braun and Clarke 2022; Naeem et al. 2023). The design integrates purposive expert sampling (Ahmad and Wilkins 2024), a standardized four-domain conversation protocol, and inductive thematic analysis following Braun and Clarke's (2022) updated reflexive approach.

3.1 Research Design Justification and Rigor

Three arguments support the research design. First, the research question requires inductive theory-building from practitioner experience. The three barrier categories identified - ontological deficit, knowledge infrastructure deficit, and decision context deficit did not appear in prior literature and could not have been hypothesized in advance. Structured surveys measure known constructs; they do not discover new categories (Braun and Clarke 2022).

Second, purposive expert sampling - the deliberate selection of participants based on domain expertise and decision-making accountability is the appropriate strategy for research requiring access to information-rich organizational informants (Ahmad and Wilkins 2024). Bouncken et al. (2025) demonstrate that in management research, purposeful informant selection with systematic saturation monitoring is the standard approach for generating transferable conceptual frameworks.

Third, accessing C-suite and director-level supply chain practitioners requires network-based purposive sampling; this population rarely responds to cold-contact academic recruitment. The failure accounts gathered in this study - including remediation cost estimates, inter-agent coordination failures, and rule accumulation patterns are the kind of operationally specific disclosures that professional network access reliably generates.

Trustworthiness is established through four criteria identified in contemporary qualitative research methodology (Ahmed 2024; Braun and Clarke 2022). Table 1 maps each criterion to the specific design mechanisms deployed in this study.

Table 1. Trustworthiness Criteria and Design Responses (Ahmed 2024; Braun and Clarke 2022)

Criterion	Definition	Design Response in This Study
Credibility	Findings accurately represent participant experiences	Standardized four-domain protocol; independent intercoder review ($\kappa = 0.79$); member checking with 6 participants; two anticipated themes excluded when prevalence was insufficient
Transferability	Findings are applicable to other contexts	Purposive sampling across 4 sectors, 3 seniority levels, $500M–$50B+ organizations; thick practitioner profiles in Appendix B
Dependability	Findings would be consistent if replicated	Fixed protocol; standardized post-engagement documentation template; transcription within 24 hours; six-phase Braun and Clarke (2022) analysis with independent coding
Confirmability	Findings are grounded in data, not researcher bias	Inductive emergence documented; non-emerging themes reported explicitly; positionality statement with analytical risk acknowledgment

3.2 Sample

Thirty-two practitioners were selected through purposive criterion sampling (Ahmad and Wilkins 2024). Sampling criteria required: (1) direct accountability for AI or technology strategy decisions in a supply chain function; (2) a minimum of five years of relevant seniority; and (3) current or recent involvement in an AI deployment decision. Participants were identified through the author's professional network, the Technology in Supply Chain Advisory Board at Saint Louis University, and the ISCEA Americas practitioner community. Sampling continued until thematic saturation was reached at approximately the 24th engagement, consistent with Bouncken et al.'s (2025) finding that saturation in focused expert samples typically occurs between 20 and 30 informants.

Participants represented organizations ranging from $500M to over $50B in annual revenue. Table 2 summarizes the sample composition.

Table 2. Sample Composition (n = 32)

Dimension	Category	n	%
Role	C-Suite / VP (CSCO, CPO, CTO, VP Supply Chain)	9	28%
	Director (Supply Chain, Technology, Operations)	13	41%
	Senior Manager (Planning, Logistics, IT)	10	31%
Industry	Retail (omni-channel, e-commerce)	11	34%
	Manufacturing (discrete and process)	9	28%
	Logistics and Third-Party Providers	7	22%
	Consumer Goods and CPG	5	16%

3.3 Data Collection Protocol
A standardized four-domain conversation protocol was developed prior to data collection and applied consistently across all 32 engagements: (1) the organization's AI investment history and stated goals; (2) specific applications attempted and their deployment outcomes; (3) the participant's own diagnosis of why adoption succeeded or stalled; and (4) what infrastructure or capability investments they believed would most accelerate adoption. The protocol was applied semi-structurally: all four domains were addressed in every engagement, but sequencing and probing depth were adjusted to the participant's experience. The protocol did not introduce the specific barriers identified in this paper — ontologies, knowledge graphs, and context structures emerged inductively from participant language, consistent with the reflexive thematic analysis approach of Braun and Clarke (2022).

Structured notes were recorded using a fixed documentation template during or immediately following each engagement and transcribed to a standardized format within 24 hours. Engagements averaged approximately 40 minutes, ranging from 20-minute focused discussions to discussions exceeding an hour.

3.4 Analysis
Analysis followed the six-phase reflexive thematic analysis process of Braun and Clarke (2022): data familiarization, initial code generation, theme construction, theme review, theme definition, and report production. Braun and Clarke's (2022) updated reflexive approach was selected because it explicitly addresses the researcher's active role in theme construction and provides clearer guidance on distinguishing themes from topics. Naeem et al. (2023) provide additional procedural guidance on developing conceptual models from thematic analysis, which informed the translation of themes into the SCALE framework.

Initial coding was conducted by the author and independently reviewed by a research collaborator with supply chain management expertise. Intercoder reliability was $\kappa = 0.79$ (Landis and Koch 1977), indicating strong agreement. Two themes anticipated prior to data collection did not emerge with sufficient prevalence to warrant inclusion: AI literacy gaps (present in six records, consistently as a secondary explanation) and vendor-capability disappointment (present in eight records, consistently attributed to infrastructure unreadiness). Their absence supports confirmability — the analysis reflects the data, not prior expectations (Braun and Clarke 2022; Ahmed 2024).

4. Findings: Three Foundational Barriers
Before examining each theme individually, the most structurally important finding warrants stating at the outset: the three barriers are not independent. They are sequential dependencies. Ontological clarity is a precondition for knowledge graph construction; a populated knowledge graph is a precondition for meaningful context accumulation; and context accumulation is a precondition for reliable multi-step agentic action. Organizations that skipped any layer and proceeded directly to agentic AI applications encountered all unresolved barriers simultaneously in production — at the point of highest remediation cost and lowest organizational patience for further investment. This sequencing pattern was present in 22 of 32 protocol records. Table 3 summarizes theme prevalence and sub-theme structure.

Table 3. Primary Themes: Prevalence and Sub-Theme Structure

Theme	Records Present	% of Sample	Primary Sub-Themes
T1: Ontological Deficit	29 of 32	91%	1a: Cross-system semantic mismatch; 1b: Absence of shared process taxonomy; 1c: Governance deterioration
T2: Knowledge Infrastructure Deficit	26 of 32	81%	2a: Tacit knowledge inaccessibility; 2b: Brittle rule-based workarounds; 2c: Temporal knowledge degradation
T3: Decision Context Deficit	24 of 32	75%	3a: Intra-cycle contradiction; 3b: Inter-agent coordination failure; 3c: Override invisibility and accountability diffusion
Cross-cutting: Sequential Dependency	22 of 32	69%	Infrastructure sequencing failures; compounding remediation costs; skipping-to-AI pattern

4.1 Theme 1: The Ontological Deficit

The most prevalent barrier — identified in 29 of 32 protocol records — was not the capability of AI models but the absence of a shared, machine-readable definition of what supply chain entities are and how they relate to one another. Leaders described this using varied language: data governance failures, master data misalignment, conflicting taxonomies, definitional inconsistency across systems. AI agents cannot reason reliably about entities — products, locations, suppliers, orders — when those entities carry different meanings in the different systems the agents must integrate.

Sub-theme 1a: Cross-system semantic mismatch

The most commonly cited manifestation, present in 18 of 32 protocol records, involved mismatched entity definitions across enterprise systems. Leaders described AI deployments that produced accurate outputs in controlled test environments and failed in production — not because the models degraded, but because production data carried definitional inconsistencies that the test environment had obscured. The operational consequence was systematic: AI recommendations built on a misread entity definition were structural errors that persisted until resolved at the data architecture level — a process practitioners universally described as expensive and slow.

Sub-theme 1b: Absence of a shared process taxonomy

Thirteen practitioners, predominantly from manufacturing and consumer goods sectors, described the absence of a formal process ontology connecting concepts across planning, procurement, and execution. Without a shared specification of how a production order relates to a purchase requisition, a work order, and a delivery schedule, agents operating across system boundaries made decisions that were internally coherent but contextually disconnected — re-initiating actions already in progress, contradicting decisions made by adjacent systems, or failing to recognize that downstream constraints made upstream recommendations infeasible.

Sub-theme 1c: Governance deterioration

Eleven protocol records addressed not the initial absence of semantic standards but their deterioration over time. Leaders described organizations that had conducted data harmonization work prior to AI deployment and found the investment had decayed: definitions drifted as systems were updated, organizational units diverged in their usage of shared terms, and governance structures were deprioritized after initial deployment. These practitioners consistently recommended that ontology governance be treated as an ongoing operational process. Interpreted through the TAM lens, the ontological deficit is a Perceived Ease of Use (PEOU) failure at the infrastructure level: when agents operate on misaligned entity definitions, every practitioner interaction is harder than it should be regardless of interface design (Davis 1989).

4.2 Theme 2: The Knowledge Infrastructure Deficit

The second theme — identified in 26 of 32 protocol records — addressed what practitioners consistently described as the absence of relational memory or network intelligence in deployed AI systems: the absence of a persistent, queryable knowledge structure encoding the relational properties of the supply chain network that experienced human planners carry as tacit knowledge.

Sub-theme 2a: Tacit knowledge inaccessibility
Twenty-two leaders described the gap between what their best planners knew and what their AI systems could access. Planners build relational intelligence over time: which suppliers run late in specific seasons, which carriers have higher damage rates on specific lanes, which distribution centers have informal capacity constraints not reflected in master data. This intelligence is operationally consequential — distinguishing a planner with 18 months of network experience from one with three months — and entirely invisible to AI systems that reason from flat transactional records. The operational consequence was described consistently: AI recommendations were technically derivable but contextually wrong in ways immediately obvious to experienced planners. This is a Perceived Usefulness (PU) failure in the TAM sense: a system that produces recommendations practitioners cannot act on without first correcting for relational context does not enhance job performance (Davis 1989).

Sub-theme 2b: Brittle rule-based workarounds
Seventeen practitioners described organizations that had attempted to address the knowledge gap through rule engineering. The pattern was remarkably consistent: rule sets grew rapidly, maintenance burden escalated as networks evolved, documentation degraded, and the eventual outcome was an AI system governed by a combination of model logic and undocumented rule logic that no individual understood in full — a governance failure that undermined rather than supported adoption.

Sub-theme 2c: Temporal knowledge degradation
Ten protocol records addressed the need for AI systems to reason about how supply chain relationships change over time. AI systems trained on static historical datasets either ignored supply disruption discontinuities or overreacted to them, without a mechanism for the contextual, temporally-weighted updating that experienced planners perform routinely. Practitioners noted this was particularly consequential during supply disruption recovery periods — precisely when agentic AI capability would otherwise be most valuable.

4.3 Theme 3: The Decision Context Deficit
The third theme — identified in 24 of 32 protocol records — was the most difficult for practitioners to articulate precisely, because it describes an absence rather than a presence. Leaders consistently described AI systems that operated as if each decision cycle were the first: without knowledge of what the system had decided in prior cycles, without awareness of what other agents were concurrently doing, and without visibility into how human operators had been modifying outputs.

Sub-theme 3a: Intra-cycle contradiction
Identified in 20 of 32 protocol records, this was the most frequent consequence: high-volume recommendation systems generating thousands of decisions per day without awareness of which prior recommendations were pending, approved, or executing, consistently producing internally contradictory decision sets. Practitioners described having to build manual reconciliation processes to resolve these contradictions — a labor cost that directly offset the efficiency gains the AI was meant to deliver.

Sub-theme 3b: Inter-agent coordination failure
Fifteen leaders described organizations with multiple AI systems operating on overlapping supply chain domains without a shared mechanism for observing each other's decisions. The resulting coordination failures ranged from redundant actions to conflicting actions to cascading failures — typically discovered after the fact through operational exception reports rather than system alerts.

Sub-theme 3c: Override invisibility and accountability diffusion
Thirteen practitioners described a governance consequence of context deficit: the disappearance of accountability when human overrides were not systematically recorded and made visible to the system. When adverse outcomes occurred, it was impossible to determine whether they resulted from AI decisions, human overrides, or the interaction between the two. This is a Trust failure in the TAM framework: stateless agents that contradict prior decisions, conflict with adjacent systems, and produce unauditable override histories are systems that practitioners cannot trust to perform as expected — because their behavior in any given cycle is disconnected from the accumulated operational state that would make their actions predictable.

4.4 Cross-Cutting Finding: Sequential Dependency and Compounding Remediation Costs

Organizations that had invested in knowledge infrastructure without first resolving ontological alignment found their graphs encoded conflicting entity definitions requiring expensive remediation. Organizations that had built context-tracking mechanisms without a knowledge graph foundation found their context records lacked the relational richness to be actionable. The pattern most frequently described — the skipping-to-AI approach — resulted in remediation costs consistently exceeding the original AI investment. The inverse was also observed: leaders who sequenced investments deliberately described materially smoother adoption trajectories and lower investment-to-capability ratios.

Figure 1 illustrates the SCALE enabling path: a sequential infrastructure investment in which each layer enables the next, leading to reliable agentic AI operation on coherent foundations. Organizations that skip these foundational layers and deploy agentic AI directly encounter all three barriers simultaneously in production — the skipping-to-AI failure pattern identified in 22 of 32 structured conversations.

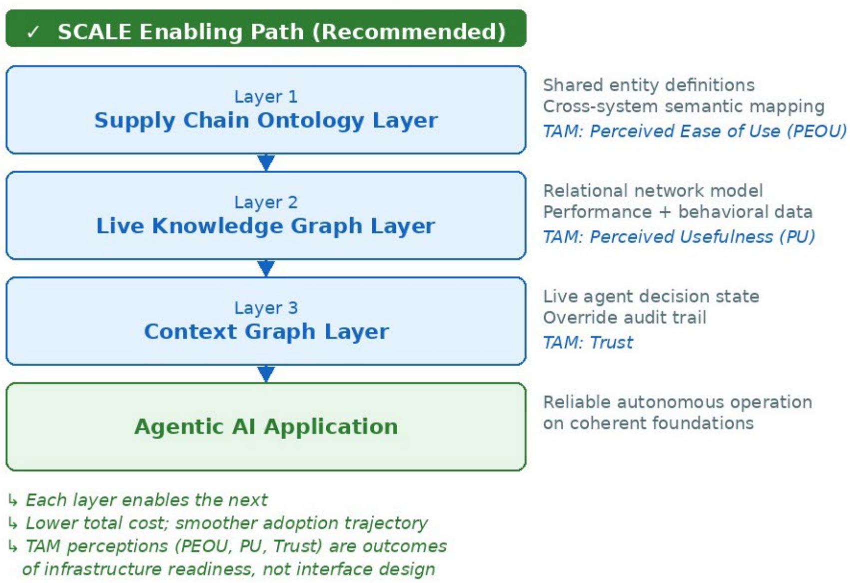

Figure 1. SCALE : The Enabling Path to Agentic AI in Supply Chain

5. Scale: Supply Chain Agentic Layer Enablement

The practitioner insights motivate a prescriptive framework organized around the three identified barriers. SCALE specifies three interdependent infrastructure layers, each addressing one barrier category. SCALE is an infrastructure blueprint, not a software product specification: it defines what each layer must accomplish; organizations choose the technologies. The three-layer dependency structure is consistent with deployment experience in a Fortune 50 retail environment, where ontological alignment across six enterprise source systems preceded knowledge graph construction, which in turn preceded the deployment of autonomous carrier assignment and replenishment logic — in that sequence, and not successfully in any other order attempted.

5.1 Layer 1 — Supply Chain Ontology Layer

The Ontology Layer provides the shared semantic foundation that agents require to reason coherently about supply chain entities across system boundaries. Its core specification is organized around five entity classes that appeared universally across practitioner conversations: Product (SKU, variant, category hierarchies), Location (stores, DCs, ports, supplier sites with typed relationships), Party (suppliers, carriers, customers, internal business units), Process (order, shipment, production run, replenishment cycle with sequencing constraints), and Constraint (capacity, lead time, compliance, contractual limits).

For each class, the ontology specifies a canonical definition, a typed relationship set, and a mapping interface aligning organization-specific terminology to canonical concepts without requiring source system changes. This reflects the key practitioner finding: the barrier is not absent data but absent interpretation layer. The governance requirement emerging from sub-theme 1c is addressed through a versioned maintenance protocol requiring change documentation, impact analysis, and domain-owner approval before deployment. Ontology governance is an ongoing operational function, not a project deliverable.

5.2 Layer 2 — Live Knowledge Graph Layer

The Knowledge Graph Layer provides agents with persistent, queryable, relational context about the supply chain network — the structural equivalent of the tacit knowledge experienced planners develop over 18 months of network exposure. The layer is built on the ontology layer's schema and populates it across three relationship categories: Structural relationships (supplier-facility links, carrier lane coverage, stocking policies — derivable from ERP and TMS master data); Performance relationships (supplier delivery rates, carrier damage rates, DC throughput — derived from historical transactions on a defined cadence); and Behavioral relationships (seasonal lead time shifts, substitution histories, override patterns by planner type — the most valuable and least accessible, requiring practitioner annotation tooling and inference pipelines from override logs).

Every relationship carries two metadata attributes: a confidence score based on evidence quality and recency, and a provenance record identifying the source and last validation date. These attributes allow agents to weight relationships appropriately and flag decisions dependent on potentially stale knowledge — the specific gap identified in sub-theme 2c.

5.3 Layer 3 — Context Graph Layer

The Context Graph Layer provides the live record of operational state practitioners consistently identified as urgently absent. Formally, the context graph C = (G, D, T, M, H) extends the knowledge graph G with four operational layers. D is the set of decision provenance nodes (agent identity, action, policy version, confidence score, timestamp). T is the temporal reasoning trace edge set, connecting sequential decision nodes. M is the inter-agent message edge set, typed by function (INFORM, REQUEST, CONFIRM, ESCALATE). H is the human override audit trail, recording practitioner interventions as typed edges (OVERRIDE, ACCEPT, DEFER, ESCALATE).

The H-layer serves a compound purpose. For practitioners, it creates a legible record of how autonomous decisions were modified, directly answering the accountability questions that sub-theme 3c identified as a recurring friction point. For agents, it is a structured feedback signal: systematic override patterns can be surfaced as behavioral relationships in the knowledge graph, closing the loop between operation and knowledge accumulation. Before generating any recommendation, an agent queries the context graph to establish what prior recommendations are active, which have been overridden and how, what other agents have decided, and what downstream effects are propagating. This transforms agents from stateless recommenders into state-aware actors capable of genuine multi-step coordination.

5.4 SCALE Integration: The Enabling Stack

Table 4 summarizes SCALE as an integrated stack, mapping each practitioner-identified barrier to its corresponding layer, the capability the layer provides, the TAM construct it addresses, and the primary organizational investment required.

Table 4. SCALE Layer-Barrier Integration Summary

Barrier (n/32 records)	SCALE Layer	Core Capability	TAM Construct Addressed	Primary Investment
T1: Ontological Deficit (29/32)	Layer 1: Ontology	Shared semantic definitions; cross-system mapping; versioned governance	Perceived Ease of Use (PEOU)	Domain modeling; data stewardship; ontology governance function
T2: Knowledge Infrastructure (26/32)	Layer 2: Knowledge Graph	Relational, queryable network model; confidence scoring; provenance	Perceived Usefulness (PU)	Graph database; extraction pipelines; practitioner annotation tooling

Barrier (n/32 records)	SCALE Layer	Core Capability	TAM Construct Addressed	Primary Investment
T3: Decision Context (24/32)	Layer 3: Context Graph	Live decision state; inter-agent coordination record; override audit trail	Trust	Append-only event store; agent API integration; override capture interface

6. Discussion

6.1 Reframing the Adoption Problem

The dominant academic framing of AI adoption barriers is individual and attitudinal: practitioners fail to adopt AI because they do not trust it, do not find it useful, or find it too complex. TAM operationalizes these as perceptual constructs and prescribes perception-level interventions. The practitioner intelligence gathered in this study suggests that PU, PEOU, and Trust are not only perceptual states — they are also outcomes of infrastructure conditions that either produce or preclude trustworthy, useful, easy-to-use system behavior. When those conditions are absent, perception-level interventions cannot close the gap.

The three barriers map directly onto TAM's core constructs. The ontological deficit produces PEOU failure: practitioners experience systems built on misaligned semantic foundations as effortful regardless of interface quality. The knowledge infrastructure deficit produces PU failure: systems without relational network knowledge generate recommendations that add verification burden rather than reduce planning effort. The decision context deficit produces Trust failure: stateless agents that contradict prior decisions and produce unauditable override histories cannot be trusted to act on behalf of practitioners who are accountable for outcomes. A planner who overrides a recommendation that contradicts a decision already in process is not exhibiting algorithm aversion — she is exhibiting an accurate perception of a system whose PU, PEOU, and Trust deficits are infrastructure-caused.

If PU, PEOU, and Trust have infrastructure preconditions, then TAM-based adoption studies in supply chain AI may be systematically misattributing infrastructure failures to individual attitudes. Future TAM instruments in supply chain AI contexts should include infrastructure readiness items as moderating variables — measuring not only whether practitioners perceive the system as useful or easy, but whether the system has the ontological, relational, and contextual foundations that make those perceptions possible.

6.2 The Sequencing Imperative

The sequential dependency finding has practical urgency. The current vendor-driven AI market presents agentic application capabilities before disclosing infrastructure requirements. Demonstrations succeed in controlled environments with clean, harmonized data; production deployments encounter the full stack of foundational deficits simultaneously. Leaders in this study who had experienced this pattern described it as one of the most costly technology investment errors they had made. SCALE makes the sequencing discipline explicit: Layer 1 before Layer 2, Layer 2 before Layer 3, and only then deployment of agentic AI applications.

6.3 Connections to Existing Literature

The findings extend Kosasih et al.'s (2023) argument about structural explainability failures in a specific direction: it is not only that AI outputs lack domain knowledge encoding, but that the foundational layers required to provide that encoding do not exist as standard enterprise infrastructure. AlMahri et al.'s (2025b) finding that knowledge graph relationship prediction enables supply chain visibility unavailable through conventional data architectures reinforces this: the infrastructure gap is real, measurable, and consequential. The sequential dependency finding connects to the path-dependency literature (Arthur 1989; David 1985): infrastructure investments create option value for subsequent application investments; skipping them forecloses options rather than merely delaying them.

6.4 Limitations

Three limitations warrant acknowledgment. First, the study sample is not random. Purposive sampling through professional networks overrepresents organizations at higher AI maturity stages; organizations at earlier stages, where foundational barriers may not yet have been encountered, are underrepresented. Second, positionality risk is present despite the intercoder review: the author's background in retail supply chain AI deployment may have introduced framing that emphasized architectural explanations. The high prevalence of Theme 1 across records is robust to this risk; the specific framing of three barriers as an ordered stack may be influenced by this perspective. Third, SCALE has not been validated through a controlled deployment study. Whether implementing SCALE reduces adoption barriers in practice requires longitudinal field research that this paper does not provide.

7. Conclusion

Thirty-two purposively sampled supply chain leaders identified, in consistent and specific terms, why agentic AI has not delivered on its promise in their organizations. Their explanations did not center on model inadequacy, practitioner resistance, or insufficient training. They centered on systems that operated without shared semantic understanding of the supply chain entities they were managing, without access to the relational intelligence that experienced planners take for granted, and without awareness of what they had already decided or what other systems were concurrently doing. These are infrastructure failures. They require infrastructure solutions.

The thematic analysis of structured protocol records produced three barriers that map cleanly to three infrastructure gaps: the absence of a shared domain ontology, the absence of a relational knowledge graph, and the absence of a live context graph. SCALE addresses each through a dedicated layer, specifies the organizational investment each layer requires, and makes explicit the sequential dependency structure that determines the order in which those investments must be made. The framework does not make AI models more capable. It creates the conditions under which capable models can function, which is the precondition the field has systematically underexamined.

TAM frames adoption as a function of perceived usefulness, perceived ease of use, and trust. This study suggests those perceptions in agentic supply chain AI contexts are shaped not primarily by interface design or change management, but by whether the system has the ontological clarity, relational knowledge, and operational context that make trustworthy, useful, and low-effort behavior possible. The path to adoption runs through the infrastructure layers that SCALE specifies.

References

All references are organized alphabetically. IEOM reference style is applied throughout.

Ahmad, M. and Wilkins, S., Purposive sampling in qualitative research: A framework for the entire journey, Quality and Quantity, 2024, https://doi.org/10.1007/s11135-024-02022-5.

Ahmed, S. K., The pillars of trustworthiness in qualitative research, Journal of Medicine, Surgery, and Public Health, vol. 2, p. 100051, 2024.

AlMahri, S., Xu, L., Mak, S. and Brintrup, A., Automating supply chain disruption monitoring via an agentic AI approach, arXiv:2601.09680, 2025a.

Arthur, W. B., Competing technologies, increasing returns, and lock-in by historical events, The Economic Journal, vol. 99, no. 394, pp. 116–131, 1989.

Atlan, Context graph vs. knowledge graph: Key differences for AI, Available: https://atlan.com, Accessed April 2026, 2024.

Bao, J., Jiang, X. and Shi, J., A formal ontology for intelligent manufacturing based on ISO/IEC and NIST standards, Computers in Industry, vol. 145, p. 103812, 2023.

Bouncken, R. B., Kraus, S. and Roig-Tierno, N., Purposeful sampling and saturation in qualitative research methodologies: recommendations and review, Review of Managerial Science, 2025, https://doi.org/10.1007/s11846-025-00881-2.

Braun, V. and Clarke, V., Thematic Analysis: A Practical Guide, Sage Publications, London, 2022.

AlMahri, S., Xu, L. and Brintrup, A., Enhancing supply chain visibility with generative AI: an exploratory case study on relationship prediction in knowledge graphs, International Journal of Production Research, 2025b, https://doi.org/10.1080/00207543.2025.2543964.

David, P. A., Clio and the economics of QWERTY, American Economic Review, vol. 75, no. 2, pp. 332–337, 1985.

Davis, F. D., Perceived usefulness, perceived ease of use, and user acceptance of information technology, MIS Quarterly, vol. 13, no. 3, pp. 319–340, 1989.

EY, Revolutionizing global supply chains with agentic AI, Available: https://www.ey.com/en_us/insights/supply-chain/revolutionizing-global-supply-chains-with-agentic-ai, Accessed April 2026, 2025.

Foundation Capital, Context is the new AI infrastructure, Available: https://reworked.co/knowledge-findability/context-is-the-new-ai-infrastructure/, Accessed April 2026, 2023.

Gartner, Gartner survey shows just 23% of supply chain organizations have a formal AI strategy, Gartner Newsroom, Available: https://www.gartner.com/en/newsroom/2025-06-11, Accessed April 2026, 2025a.

Gartner, Gartner predicts 60% of supply chain digital adoption efforts will fail to deliver promised value by 2028, Gartner Newsroom, Available: https://www.gartner.com/en/newsroom/press-releases/2025-05-07, Accessed April 2026, 2025b.

Gartner, Gartner predicts half of supply chain management solutions will include agentic AI capabilities by 2030, Gartner Newsroom, Available: https://www.gartner.com/en/newsroom/press-releases/2025-05-21, Accessed April 2026, 2025c.

Gartner, Hype cycle for supply chain strategy 2025, Gartner Research, Available: https://www.gartner.com, Accessed April 2026, 2025d.

Gruber, T. R., A translation approach to portable ontology specifications, Knowledge Acquisition, vol. 5, no. 2, pp. 199–220, 1993.

Hangl, J., Krause, S. and Behrens, V. J., Drivers, barriers and social considerations for AI adoption in supply chain management: a tertiary study, Technology in Society, vol. 74, p. 102299, 2023.

Ibrahim, F., Münscher, J.-C., Daseking, M. and Telle, N.-T., The technology acceptance model and adopter type analysis in the context of artificial intelligence, Frontiers in Artificial Intelligence, vol. 7, p. 1496518, 2025.

Jackson, I., Ivanov, D., Dolgui, A. and Namdar, J., Generative artificial intelligence in supply chain and operations management: a capability-based framework for analysis and implementation, International Journal of Production Research, vol. 62, no. 17, pp. 6120–6145, 2024.

Jannelli, V., Schoepf, S., Bickel, M., Netland, T. and Brintrup, A., Agentic LLMs in the supply chain: towards autonomous multi-agent consensus-seeking, International Journal of Production Research, 2025, https://doi.org/10.1080/00207543.2025.2604311.

Ji, S., Pan, S., Cambria, E., Marttinen, P. and Yu, P. S., A survey on knowledge graphs: representation, acquisition, and applications, IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 2, pp. 494–514, 2021.

Kluza, K., Jobczyk, K., Wisniewska, P. and Ligeza, A., Overview of selected ontologies for supply chain management, Applied Sciences, vol. 12, no. 12, p. 5903, 2022.

Kosasih, E. E. and Brintrup, A., Towards knowledge graph reasoning for supply chain risk management using graph neural networks, International Journal of Production Research, vol. 60, no. 17, pp. 5380–5393, 2022.

Kosasih, E. E. and Brintrup, A., Towards trustworthy AI for link prediction in supply chain knowledge graph: a neurosymbolic reasoning approach, International Journal of Production Research, vol. 62, no. 24, pp. 2268–2290, 2024.

Kosasih, E. E., Srinivasan, R. and Brintrup, A., A review of explainable artificial intelligence in supply chain management using neurosymbolic approaches, International Journal of Production Research, vol. 61, no. 11, pp. 3658–3674, 2023.

Landis, J. R. and Koch, G. G., The measurement of observer agreement for categorical data, Biometrics, vol. 33, no. 1, pp. 159–174, 1977.

Li, B., Mellou, K., Zhang, B., Pathuri, J. and Menache, I., Large language models for supply chain optimization, arXiv:2307.03875, 2023.

Naeem, M., Ozuem, W., Howell, K. and Ranfagni, S., A step-by-step process of thematic analysis to develop a conceptual model in qualitative research, International Journal of Qualitative Methods, vol. 22, 2023.

Noy, N. F. and McGuinness, D. L., Ontology development 101: a guide to creating your first ontology, Stanford Knowledge Systems Laboratory Technical Report KSL-01-05, 2001.

Queiroz, M. M., Ivanov, D., Fosso Wamba, S. and Wanke, P., Perceptions of AI adoption and their impact on supply chain learning, International Journal of Logistics Research and Applications, 2025, https://doi.org/10.1080/13675567.2025.2547202.

SAP, Agentic AI in the global supply chain, SAP Insights, Available: https://www.sap.com/blogs/agentic-ai-in-global-supply-chain, Accessed April 2026, 2025.

Shahzadi, I., Arshad, M. U., Farooq, M. U. and Abbas, M., AI adoption in supply chain management: a systematic literature review, ResearchGate Preprint, 2024, https://doi.org/10.13140/RG.2.2.34567.89045.

World Economic Forum, Autonomous orchestration: the next frontier in supply chain management, WEF Insight Report, Available: https://www.weforum.org/stories/2025/11/autonomous-orchestration-next-frontier-supply-chain-management/, Accessed April 2026, 2025.

Xu, L., Almahri, S., Mak, S. and Brintrup, A., Multi-agent systems and foundation models enable autonomous supply chains: opportunities and challenges, IFAC-PapersOnLine, vol. 58, no. 19, pp. 892–897, 2024.

Xu, L., Mak, S. and Brintrup, A., Will supply chain management survive the digital revolution? International Journal of Production Economics, vol. 239, p. 108171, 2021.

Xu, L. and Brintrup, A., Enhancing supply chain visibility with knowledge graphs and large language models, International Journal of Production Research, vol. 63, no. 11, pp. 2178–2209, 2025.
Zero100, Supply chain's AI building blocks: knowledge graphs vs ontologies, Zero100 Intelligence Report, Available: https://zero100.com, Accessed April 2026, 2025.

Biographies

Prabhat Rao Pinnaka is Lead Product Manager for Supply Chain Technology at Lowe's Companies, Inc. (Fortune 50). Prior to Lowe's, he led supply chain and operations consulting at Ernst & Young and Deloitte, serving Fortune 500 organizations including Kraft Heinz, Dow Chemicals, BASF, Molson Coors, and Johnson Controls. He serves on the ISCEA Americas Advisory Board and is a founding member of the Technology in Supply Chain Advisory Board at Saint Louis University. He spoke at Manifest 2026 on agentic supply chain systems. His research has appeared in Supply Chain Management Review, ISE Magazine, Automated Warehouse, DSG Magazine, and Supply & Demand Chain Executive . His work on agentic AI in supply chain systems is presented at the IEOM Milwaukee 2026 conference. He holds an M.S. in Integrated Manufacturing Systems Engineering from North Carolina State University and a B.Tech in Mechanical Engineering from NIT Warangal.

Sukanya Bollineni is a Technical Product Owner. With a background in technology and product delivery, she works at the intersection of business needs, digital solutions, and cross-functional execution. Her professional focus aligns with broader efforts in digital manufacturing and connected industrial operations, where technology is increasingly used to improve efficiency, performance, and smarter decision-making.

Appendix A: Conversation Guide

The following topic areas and illustrative prompts structured the practitioner conversations. Prompts were used as discussion anchors and applied semi-structurally across all 32 engagements. The language below reflects the actual framing used, not post-hoc reconstruction.

A.1 Opening Orientation (2–5 minutes)

Each conversation opened with a brief orientation to establish the scope of the practitioner's accountability and their organization's stage of AI deployment.

Illustrative prompts: "How would you describe your organization's current position on AI in supply chain — are you evaluating, piloting, or operating deployed systems?" / "Where does your accountability sit in that journey — are you the one making the architecture decisions, the deployment decisions, or both?"

A.2 AI Investment and Stated Goals

This section probed the nature and scale of AI investments, the supply chain functions targeted, and the capability outcomes intended.

Illustrative prompts: "What supply chain decisions were you trying to automate or augment, and what level of autonomy were you aiming for?" / "How would you describe the gap between what you were promised and what you got?"

A.3 Deployment Outcomes and Failure Points

This section asked participants to describe what happened when AI deployments moved from pilot to production, with attention to specific failure points.

Illustrative prompts: "Walk me through what happened when you took this from pilot to live. Where did it start behaving differently than you expected?" / "If you had to point to the single decision or gap that most explains why the system didn't perform as expected in production, what would it be?"

A.4 Root Cause Diagnosis

This section asked participants to offer a diagnostic judgment about the category of problem they had encountered.

Illustrative prompts: "If someone asked you to explain in one sentence why AI is harder to deploy in supply chain than the vendors suggest, what would you say?" / "If the AI model itself were perfect — if it always produced the right recommendation — would your organization be able to act on it reliably? Why or why not?"

A.5 Infrastructure and Enablement

The closing section asked what foundational investments would most accelerate adoption. This section produced the richest material on ontological, knowledge, and context infrastructure needs.

Illustrative prompts: "If you were starting over with a clean slate, what would you build first, before you touched an AI model?" / "What does the AI need to know about your supply chain that it currently doesn't have access to — and why isn't that information available to it?" / "What would have to be true about your data and systems infrastructure for you to feel comfortable letting the AI act autonomously?"

Appendix B: Anonymized Practitioner Profiles

Table B1 provides anonymized profiles of all 32 practitioners whose conversations informed the thematic analysis. All identifying details — names, organization names, and precise revenue figures — have been removed or generalized. Role titles reflect functional accountability. C-suite leaders from technology solution provider startups are identified by functional role only to prevent identification.

Table B1. Retail Sector Practitioners (n = 11)

ID	Role	Organization Type	Revenue Range	Key Contribution to Analysis
P-R01	VP Supply Chain Technology	Large-format omni-channel retailer	$15B–25B	AI and automation strategy across fulfillment; second year of autonomous replenishment pilot
P-R02	VP Supply Chain Operations	National specialty retailer	$5B–10B	Attempted carrier assignment automation; deployment rolled back after six months
P-R03	Director, Inventory Planning	Omni-channel retailer	$10B–20B	Primary informant on intra-cycle contradiction sub-theme (T3a)
P-R04	Director, Supply Chain Technology	E-commerce retailer	$2B–5B	Led AI vendor evaluation; organization had not proceeded to deployment
P-R05	Director, Logistics Technology	Omni-channel retailer	$8B–15B	Carrier integration and routing optimization; high planner override rate
P-R06	Director, Digital Supply Chain	National grocery retailer	$3B–6B	Described ontological deficit as primary barrier to AI integration across ERP and WMS
P-R07	Senior Manager, Replenishment Technology	Large-format retailer	$20B+	Autonomous replenishment pilot lead; primary informant on rule workaround sub-theme (T2b)
P-R08	Senior Manager, Supply Chain Analytics	Specialty retailer	$1B–2B	Organization at early AI evaluation stage; focused on infrastructure readiness
P-R09	Senior Manager, Fulfillment Operations	E-commerce retailer	$4B–8B	Described inter-agent coordination failure involving two separately deployed AI systems
P-R10	Senior Manager, Logistics Planning	Omni-channel retailer	$6B–12B	Primary informant on temporal knowledge degradation sub-theme (T2c)
P-R11	Senior Manager, Inventory Technology	Discount retailer	$2B–4B	Described semantic mismatch between inventory and order management systems

Table B2. Manufacturing Sector Practitioners (n = 9)

ID	Role	Organization Type	Revenue Range	Key Contribution to Analysis
P-M01	Chief Supply Chain Officer	Industrial manufacturer	$8B–18B	Primary informant on sequencing meta-finding; described skipping-to-AI pattern directly
P-M02	VP Digital Supply Chain	Discrete manufacturer	$5B–12B	Built ontological alignment layer prior to AI deployment; primary enabler of adoption success
P-M03	VP Procurement Technology	Process manufacturer	$10B–25B	Absence of supplier knowledge graph as binding constraint on autonomous procurement
P-M04	Director, Manufacturing Planning	Discrete manufacturer	$3B–7B	Primary informant on process taxonomy absence sub-theme (T1b)
P-M05	Director, Supply Chain Technology	Consumer electronics manufacturer	$4B–8B	Semantic standardization effort that decayed over 18 months; sub-theme (T1c)
P-M06	Director, Operations Technology	Specialty chemicals manufacturer	$2B–5B	AI scheduling system requiring 340 business rules to operate acceptably
P-M07	Senior Manager, Supply Chain Analytics	Automotive components manufacturer	$1B–3B	18-month planner learning curve as benchmark for AI knowledge gap; sub-theme (T2a)
P-M08	Senior Manager, Digital Manufacturing	Industrial equipment manufacturer	$3B–6B	Cross-system semantic mismatch in manufacturing-procurement integration
P-M09	Senior Manager, Procurement Operations	Process manufacturer	$2B–4B	Brittle rule workaround; over 200 procurement-specific rules accumulated in two years

Table B3. Logistics and Third-Party Provider Practitioners (n = 7)

ID	Role	Organization Type	Revenue Range	Key Contribution to Analysis
P-L01	VP Logistics Technology	Third-party logistics provider	$5B–12B	Primary informant on inter-agent coordination failure sub-theme (T3b)
P-L02	VP Technology	Regional 3PL	$800M–2B	AI routing producing internally contradictory recommendations at scale; intra-cycle contradiction (T3a)
P-L03	Director, Network Optimization	National freight carrier	$3B–7B	Override invisibility as persistent accountability gap in commercial leadership (T3c)
P-L04	Director, Logistics Technology	Technology solution provider (startup)	N/A	Ontological deficit from the vendor side: agentic logistics AI when customer data models are not standardized
P-L05	Director, Operations Technology	Third-party fulfillment provider	$1B–2B	Primary informant on distinction between structural and behavioral knowledge relationships
P-L06	Senior Manager, Data and Analytics	Freight brokerage	$500M–1B	Semantic mismatch between internal and carrier-provided location data causing routing failure

ID	Role	Organization Type	Revenue Range	Key Contribution to Analysis
P-L07	Senior Manager, Technology Solutions	Technology solution provider (startup)	N/A	Vendor-side perspective on foundational infrastructure gaps encountered during customer deployments

Table B4. Consumer Goods and CPG Practitioners (n = 5)

ID	Role	Organization Type	Revenue Range	Key Contribution to Analysis
P-C01	VP Supply Chain	Multinational CPG company	\$10B–20B	Skipping-to-AI pattern; remediation cost estimated at 60% above original AI investment
P-C02	Director, Demand Planning Technology	Mid-size CPG company	\$2B–5B	Temporal knowledge degradation during post-COVID demand volatility; sub-theme (T2c)
P-C03	Director, Procurement	Large CPG company	\$8B–15B	Tacit knowledge inaccessibility in procurement; seasonal supplier behavior not encoded in any system
P-C04	Senior Manager, Supply Chain Technology	Regional CPG company	\$500M–1B	Organization not yet attempted agentic AI; focused on infrastructure readiness assessment
P-C05	Senior Manager, Logistics and Distribution	Technology solution provider (startup)	N/A	Knowledge infrastructure requirements for autonomous inventory positioning; prior Fortune 100 CPG experience