NextGen Scientific Publication
American International Journal of Computer Science and Technology

Volume 8 Issue 6, Pg. No. 16-21, AIJCST-V8I6P102, 2026
https://doi.org/10.63282/3117-5481/AIJCST-V8I6P102

Original Article

Next-Generation Autonomous Procurement: An Agentic Multi-Agent Architecture for the Source-to-Procure Workflow

***Deepika Nathany**
Application Manager, USA.

Abstract:

The evolution of Enterprise Resource Planning (ERP) and Supply Chain Management (SCM) systems has reached a pivotal juncture with the transition from traditional, rule-based automation to agentic artificial intelligence (AI). Traditional Source-to-Procure (S2P) workflows remain constrained by manual human handoffs, static threshold approvals, and systemic silos spanning Product Lifecycle Management (PLM), demand planning, sourcing, and ERP purchasing systems. This paper introduces a multi-agent architectural framework for an agentic Source-to-Procure workflow. By leveraging specialized, autonomous AI agents—a PLM-focused Bill of Materials (BOM) Validator Agent, a Demand Planning Agent, a Sourcing and Optimization Agent, and a Purchase Order (PO) Execution Agent—under a centralized orchestration layer, the proposed system operates via cross-functional context handoffs. We analyze the technical architecture, tool protocols, security boundaries, and industry trends driving this evolution, detailing how agentic S2P substantially reduces cycle times, improves procurement accuracy, and enhances enterprise resilience.

Keywords:

Agentic AI, Autonomous Procurement, Source-To-Procure (S2P), Multi-Agent System (MAS), Bill of Materials (BOM), Orchestration Layer, Enterprise Security.

Article History:

Received: 18.09.2026

Revised: 16.10.2026

Accepted: 26.10.2026

Published: 04.11.2026

1. Introduction

Global supply chains face unprecedented volatility, compressed product lifecycles, and increasingly fragmented data ecosystems across multi-tiered supplier networks. Historically, artificial intelligence in supply chain management has focused primarily on predictive analytics, such as forecasting demand patterns, assessing supply risk, or optimizing route logistics through machine learning (ML) models (Baryannis et al., 2019; Boute et al., 2021; Min, 2010). While these models deliver actionable insights, the execution of downstream actions—validating part specifications, verifying demand forecasts, negotiating sourcing options, and issuing purchasing contracts—has remained heavily reliant on manual labor and static ERP workflows.

The emergence of agentic AI marks a paradigm shift from insight generation to autonomous task execution. Unlike deterministic algorithms or legacy Robotic Process Automation (RPA) that depend on rigid, predefined scripts, AI agents operate with goal-directed behavior, dynamic reasoning, and execution capabilities over domain-specific tools.

Applying agentic architectures to the Source-to-Procure (S2P) process enables enterprises to bridge the gap between engineering (PLM), supply planning (Sales and Operations Planning [S&OP] and demand planning), procurement strategy, and ERP order

Copyright by the Author(s). This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-NC-ND 4.0) License (https://creativecommons.org/licenses/by-sa/4.0/)

generation. This paper introduces a modular, four-agent orchestration architecture for the S2P lifecycle, establishing security boundaries, tool interaction protocols, and system dynamics to model the future of autonomous, resilient enterprise procurement.

2. Industry Trends and Background

Recent industry shifts indicate a rapid acceleration from experimental generative AI pilots toward autonomous agent orchestration across enterprise functions. Key trends reshaping the S2P landscape include the following:

- Shift from assistive to agentic procurement. Enterprise applications are embedding autonomous agents directly into operational workflows, moving from static dashboards to zero-click intelligence and action-oriented execution.
- Adoption of standardization protocols. The emergence of standardized frameworks such as the Model Context Protocol (MCP) allows AI agents to interface natively with enterprise software databases, vendor portals, and external market feeds.
- Hyper-personalized and dynamic supply networks. Continuous geopolitical shifts and regulatory pressures (e.g., environmental, social, and governance [ESG] and carbon compliance tracking) require continuous, real-time risk surveillance and supplier discovery rather than static quarterly audits.

Table 1 situates these trends within the broader evolution of procurement technology. Prior research on AI in supply chains has concentrated largely on the predictive and decision-support capabilities of the third wave (Baryannis et al., 2019; Min, 2010); the present work addresses the fourth.

Table 1. Evolution of Procurement Technology toward Agentic Source-to-Procure

Wave / Era	Core Technologies	Primary Capability	S2P Workflow Role
Procurement 1.0	Manual records and EDI	Paper-based and basic digital ordering	Human-driven data entry and static POs
Procurement 2.0	ERP systems and deterministic RPA	Rule-based workflow automation	Scripted approvals and static linear routing
Procurement 3.0	Machine learning and predictive analytics	Demand and risk forecasting	Decision-support analytics for human buyers
Procurement 4.0 (Agentic)	Multi-agent systems and tool orchestration	Autonomous, context-driven execution	End-to-end multi-agent autonomous S2P

Note: *EDI = electronic data interchange; ERP = enterprise resource planning; RPA = robotic process automation; PO = purchase order; S2P = Source-to-Procure.*

3. System Architecture and Multi-Agent S2P Workflow

The core architecture relies on a specialized multi-agent system (MAS) coordinated through an orchestration and context layer. Rather than a single large language model handling the entire workflow, domain-specific tasks are decoupled into four autonomous agents, as shown in Figure 1.

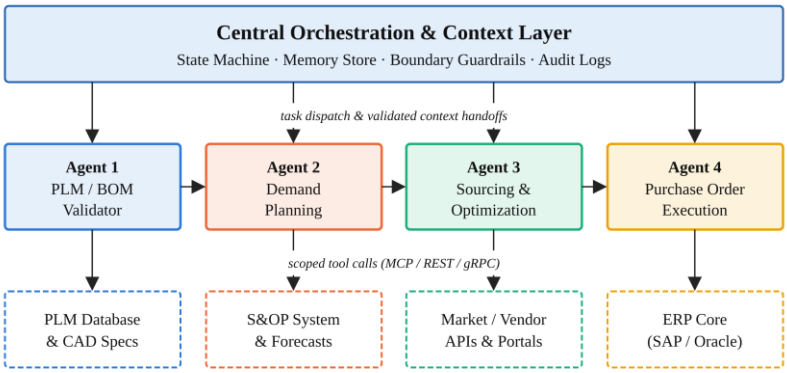


Figure 1. Multi-Agent Source-to-Procure Architecture

Note: *Solid boxes denote autonomous agents; dashed boxes denote the enterprise systems each agent accesses through scoped tools. Horizontal arrows indicate the sequential context handoffs between agents.*

3.1. Agent Roles and Functional Responsibilities

3.1.1. Agent 1: PLM BOM Validator Agent

- Primary objective: Validates engineering change orders (ECOs), raw Bill of Materials (BOM) files, component revisions, and technical specifications.
- Functional scope: Checks part numbers, verifies active lifecycle status in the PLM repository, ensures single-source consistency across sub-assemblies, and flags obsolete or end-of-life (EOL) components.

Context handover: Produces a normalized, fully validated *Clean Sourcing BOM* payload containing precise technical attributes, tolerance specifications, and the approved manufacturer part list (AMPL).

3.1.2. Agent 2: Demand Planning Agent

- Primary objective: Cross-references the Clean Sourcing BOM against master production schedules (MPS), enterprise inventory levels, and consumption forecasts.
- Functional scope: Queries the supply planning engine to determine target lead times, safety stock requirements, existing stock on hand, and open order pipelines across regional manufacturing plants.
- Context handover: Generates a *Net Purchase Requirement* payload indicating precise component quantities, target delivery dates, and plant-level receiving parameters.

3.1.3. Agent 3: Sourcing and Optimization Agent

- Primary objective: Identifies, evaluates, and selects optimal supplier options based on cost, quality, lead time, supplier capacity, and risk metrics.
- Functional scope: Accesses real-time vendor catalogs, contractual rate cards, supplier performance histories, and live risk metrics, and evaluates multi-objective optimization criteria (e.g., total cost vs. delivery speed vs. ESG score).
- Context handover: Delivers an *Approved Sourcing Recommendation* payload specifying selected vendor IDs, tiered pricing structures, terms, and delivery schedules.

3.1.4. Agent 4: Purchase Order Execution Agent

- Primary objective: Converts the sourcing recommendation into a binding purchase order (PO) and submits it to the transactional core.
- Functional scope: Interfaces directly with the enterprise ERP system via secure APIs, validates PO lines against enterprise governance parameters, attaches relevant specification files, generates contract documentation, and executes submission to the supplier portal.
- Context handover: Emits a *PO Execution Artifact* payload containing the ERP tracking ID, confirmation tokens, and audit logs to the central orchestration memory.

3.2. Agent–System Access Boundaries

Taken together, the responsibilities in Section 3.1 define a least-privilege access pattern in which each agent touches only the enterprise system its task requires. Figure 2 consolidates this pattern as an access matrix. Agents 1 through 3 hold read-only access to a single system each, whereas write privileges are concentrated exclusively in the PO Execution Agent, the only agent that commits binding transactions. This concentration of write access is what allows the security controls described in Section 4.3 to be applied in proportion to transactional risk.

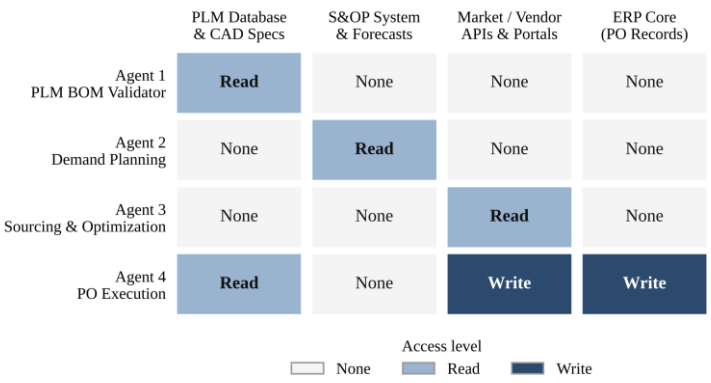


Figure 2. Agent–System Access Matrix under Least-Privilege Scoping

Note: *Access levels are derived from the agent responsibilities in Section 3.1 and the role-based access control rules in Section 4.3. Agent 4 holds read access to the PLM system to attach specification files and write access to the supplier portal and the ERP core to submit and record purchase orders.*

4. Technical Foundations: Orchestration, Protocols, and Boundaries

4.1. Orchestration Layer and Context Handoffs

The orchestration layer manages state transitions, long-term memory registers, and inter-agent communication. Using state-graph or directed acyclic graph (DAG) execution models, the orchestrator enforces strict contract schemas (e.g., JSON Schema or Pydantic models) for memory passing between agents:

- Context enforcement: Agent N cannot execute until Agent $N-1$ produces a cryptographically verified and validated state output.
- Rollback and exception handling: If Agent 3 discovers zero active suppliers capable of meeting lead times, the orchestrator triggers a fallback loop to Agent 1 to check for alternative qualified component substitutes within the PLM database.

4.2. Tool Protocols and Integration Architecture

Agents interface with enterprise environments using structured tool-use mechanisms (e.g., the Model Context Protocol or REST/gRPC OpenAPI endpoints):

- Read-only data retrieval tools: Query PLM tables, fetch S&OP demand tables, and read vendor rate cards.
- Analytical tools: Run constraint-based mathematical solvers for multi-criteria sourcing optimization.
- Write/transactional tools: Commit purchase requisitions, update inventory allocations, and post ERP PO records.

4.3. Security Framework and Boundary Conditions

Autonomy in financial transaction workflows introduces significant risk vectors, requiring robust enterprise guardrails. Figure 3 summarizes the four boundary layers, each of which is described below.

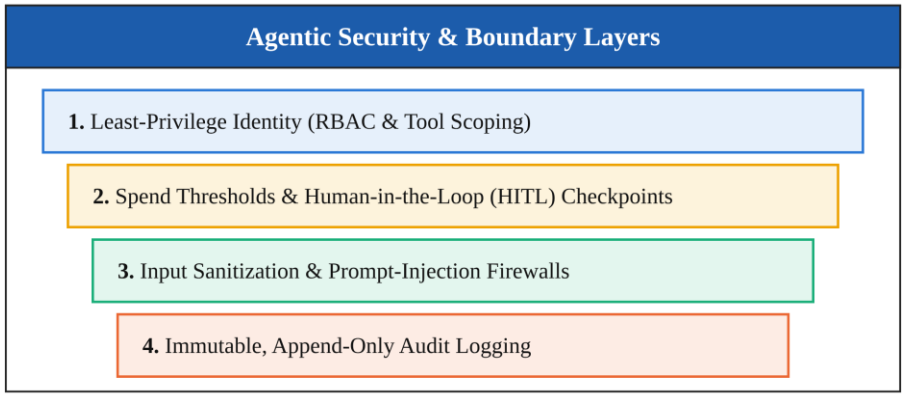


Figure 3. Layered Security and Boundary Controls for Agentic S2P

- Spend boundary thresholds (HITL checkpoints). Fully autonomous execution is governed by financial risk matrix rules. Low-value, high-frequency component orders below a specified monetary threshold (T_{spend}) execute autonomously. Purchases exceeding T_{spend} trigger a human-in-the-loop (HITL) approval gate, presenting the human buyer with an agent-generated rationale and a structured comparative choice matrix.
- Tool scoping and RBAC. Each agent operates under scoped API service credentials governed by role-based access control (RBAC). The PLM Validator Agent lacks write access to ERP financial records, and the PO Execution Agent cannot modify engineering BOMs.
- Adversarial and injection defense. External vendor inputs (e.g., vendor quotes, external product specifications) pass through an isolation sandbox and prompt-injection firewall to prevent malicious string execution within downstream agent reasoning loops.
- Immutable audit trails. Every reasoning step, tool call, external payload, and human interaction is recorded in an append-only transaction ledger for governance, compliance auditing, and post hoc process optimization, consistent with blockchain-based traceability approaches in supply chain management (Rejeb et al., 2019).

These controls are not applied uniformly. As Figure 4 shows, the number of controls governing each agent increases along the workflow, from two for the internally facing PLM and demand planning agents to four for the PO Execution Agent. Least-privilege scoping and audit logging apply to every agent; the prompt-injection firewall applies only where external vendor inputs enter the workflow (Agents 3 and 4); and the human-in-the-loop spend threshold applies only at the point of financial commitment (Agent 4). The resulting gradient aligns the strength of governance with each agent's exposure to external data and financial risk.

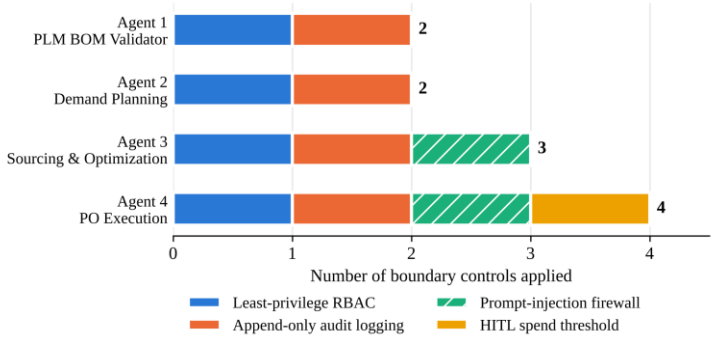


Figure 4. Boundary Controls Applied to Each Agent

Note: *Counts reflect the assignment of the four boundary layers in Figure 3 to each agent according to the framework design; they are not empirical measurements. HITL = human-in-the-loop; RBAC = role-based access control.*

5. Strategic Benefits and Business Impact

The transition from traditional S2P workflows to an agentic multi-agent architecture delivers tangible improvements across key operational performance indicators, consistent with empirical evidence linking big data analytics and AI capabilities to operational performance in manufacturing organizations (Dubey et al., 2020):

- Compression of cycle times: Reductions in human handoff latency turn traditional S2P cycles—often requiring days or weeks across PLM, procurement, and planning teams—into dynamic, near-real-time executions.
- Elimination of transposition and data errors: Autonomous BOM validation prevents costly manufacturing discrepancies, scrap, and line shutdowns caused by ordering legacy or out-of-specification revisions.
- Dynamic cost and tail-spend optimization: Sourcing agents continuously evaluate real-time rate cards, spot-market opportunities, and tiered discounts across tail spend that often escapes manual strategic sourcing oversight.
- Resilience to supply disruptions: Autonomous orchestration enables real-time, automated re-sourcing triggers in response to live supply disruption signals or sudden demand spikes, extending the visibility and flexibility benefits associated with supply chain analytics (Srinivasan & Swink, 2018).

6. Implementation Challenges and Ethical Considerations

Despite its transformative potential, deploying an agentic S2P framework introduces organizational and technical challenges:

- Legacy ERP data silos and data quality. Agentic reasoning depends on high-quality, standardized master data across PLM, S&OP, and ERP systems. Fragmented data sets impede tool-invocation accuracy, a dependence on integrated, high-quality data that the smart manufacturing literature likewise emphasizes (Kusiak, 2018).
- Governance and accountability. Determining clear operational liability for autonomous vendor commitments requires explicit organizational policy and legal frameworks.
- Change management. Transitioning procurement teams from tactical transactional processing to higher-level agent architecture governance requires structured upskilling.

7. Conclusion and Future Directions

An agentic multi-agent approach transforms the Source-to-Procure workflow into an autonomous, proactive, and resilient operational driver. By coordinating specialized agents for PLM BOM validation, demand planning, dynamic sourcing, and purchase order execution under strict security boundaries and orchestration rules, organizations can streamline S2P operations, minimize operational friction, and maintain adaptivity in uncertain global markets.

7.1. Future Research Directions

Several avenues merit further investigation:

- Implementation of decentralized multi-agent negotiation frameworks for automated, dynamic price discovery with vendor-side agents.
- Integration of real-time, multi-tier ESG and carbon-tracking metrics into agent decision loops.
- Formal validation of explainable AI (XAI) frameworks tailored to enterprise procurement audit compliance.

References

[1] Baryannis, G., Validi, S., Dani, S., & Antoniou, G. (2019). Supply chain risk management and artificial intelligence: State of the art and future research directions. *International Journal of Production Research, 57*(7), 2179–2202.

[2] Boute, R., Gijsbrechts, J., & Van Mieghem, J. A. (2021). AI in logistics and supply chain management. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3860431

[3] Dubey, R., Gunasekaran, A., Childe, S. J., Bryde, D. J., Giannakis, M., Foropon, C., . . . Fosso Wamba, S. (2020). Big data analytics and artificial intelligence pathway to operational performance under the effects of entrepreneurial orientation and environmental dynamism: A study of manufacturing organisations. *International Journal of Production Economics, 226*, Article 107599.

[4] Kusiak, A. (2018). Smart manufacturing. *International Journal of Production Research, 56*(1–2), 508–517.

[5] Min, H. (2010). Artificial intelligence in supply chain management: Theory and applications. *International Journal of Logistics Research and Applications, 13*(1), 13–39.

[6] Rejeb, A., Keogh, J. G., & Treiblmaier, H. (2019). Leveraging the Internet of Things and blockchain technology in supply chain management. *Future Internet, 11*(7), Article 161.

[7] Srinivasan, R., & Swink, M. (2018). An investigation of visibility and flexibility as complements to supply chain analytics: An organizational information processing theory perspective. *Production and Operations Management, 27*(10), 1849–1867.