

PROCUREMENT · RUNTIME GOVERNANCE

Agentic AI Is Entering *Procurement*

Approval Controls Need to Govern the Meaning Behind the Spend

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

Agentic AI Is Entering *Procurement*

Approval Controls Need to Govern the Meaning Behind the Spend

MAUREEN DOYLE-SPARE

As procurement agents move from drafting into sourcing and execution, the control question shifts from whether a transaction is allowed to whether the interpretation behind it was authorized. Governing that shift calls for runtime oversight and a deterministic decision boundary before execution.

Procurement is moving from assistance to execution

The current procurement discussion is no longer about whether generative AI can summarize a contract or draft an RFP. Procurement Magazine's August 23 coverage[1] describes a shift toward systems that can execute broader procurement workflows, from supplier evaluation and sourcing to requisitions, approvals, purchase orders and payments.

That direction is consistent with Gartner's 2026 supply-chain forecast[2], which projects spending on supply-chain software with agentic AI capabilities rising from under $2 billion in 2025 to $53 billion by 2030. Gartner also expects organizations to move toward agent orchestration across multi-step workflows, with or without humans in the loop. Procurement sits directly in that path because much of the function consists of decisions that appear rule-based until the surrounding context changes.

Its attraction is easy to understand. Procurement agents can collect requirements, identify suppliers, compare bids, prepare purchase requests, check contract terms, coordinate approvals and eventually initiate transactions. Each step can reduce manual effort and cycle time. But as the agent assumes more of that sequence, a distinction that mattered less for conventional automation becomes central: technical permission to advance a workflow is not the same thing as institutional authority for the interpretation that caused it to advance.

An agent can execute flawless steps against a meaning no organization authorized.

Approval controls begin after some of the most important decisions have already been made

Procurement functions already operate with extensive

© 2026 Maureen Doyle-Spare.

controls. Purchases may have to use an approved supplier, remain within delegated authority, satisfy category rules, meet security or quality requirements, use accepted contractual terms and receive specified approvals. Those controls do not disappear when agents enter the process.

The difficulty is that an autonomous system may have to determine which of those controls apply before it can use them: whether this supplier is approved for this particular service, geography or data environment; whether a contractual change is routine or material; whether a purchase remains within its original mandate once scope shifts; and whether an acceptable supplier-risk rating stays acceptable once the service becomes operationally critical.

Answers may depend on procurement policy, contract language, supplier-risk records, cybersecurity requirements, quality standards, budget rules, business-unit exceptions, resilience requirements and prior approved practice. Each source can be valid while the meaning produced from their combination is not.

That is the control problem agentic procurement exposes. A workflow can satisfy every downstream approval rule presented to it while the interpretation that determined the route through those rules has already moved outside the organization's authorized position.

The control surface is the interpretation before the transaction

Traditional procurement systems are good at enforcing deterministic conditions once the relevant facts and classifications have been established. Agentic systems complicate that sequence because they can participate in establishing the classification itself. An agent does more than check whether an approved-supplier flag is present; it can decide what "approved" means for the combination of service, data, location, contractual terms and operating dependency in front of it.

That decision needs to become governable before the agent can turn it into action. An organization needs an authoritative reference for the meaning that applies, a representation of the interpretation the agent has resolved, and an enforceable way to stop execution when the two diverge beyond the organization's tolerance.

The Semantic Control Plane[3] provides one architecture for that boundary. The organization sets a Reasoning Baseline for the decision class; the Semantic Deviation Index[4] measures how far the agent's interpretation has moved from it; and the Deterministic Gate returns permit, permit-with-flag, hold, or mandatory hold before execution authority is emitted.

In procurement, this does not replace sourcing policy, supplier controls, contract management, budget authority, cybersecurity review, quality requirements or human escalation. It addresses a different problem: whether the meaning entering those controls is one the organization actually authorized.

The same procurement problem changes shape across industries

A manufacturer sourcing a component may have to reconcile price and availability with quality certification, country-of-origin restrictions and approved-substitution rules. In pharmaceuticals, the question is whether a supplier fits within validated processes and quality agreements. Hospitals add clinical, privacy and cybersecurity conditions. An energy or telecommunications operator must weigh critical-infrastructure dependencies. For a retailer, it is inventory, supplier performance, margin and continuity.

Vocabulary changes, but the architecture of the decision does not. Multiple authoritative sources contribute

© 2026 Maureen Doyle-Spare.



Figure 1. *The Pre-Execution Assurance Protocol. Source: Doyle-Spare (2026)* [3].

meaning to a pending action; the agent resolves them into an interpretation; that interpretation determines what happens next. As long as a human performs most of that reconciliation, organizations can rely heavily on delegated judgment and escalation. As agents assume the reconciliation themselves, the interpretation becomes an explicit control surface.

That makes procurement a useful early test for enterprise agentic governance. It crosses organizational boundaries, connects internal policy to external counterparties, and contains enough exceptions that simply giving an agent access to the right systems does not establish that it will resolve their combined meaning correctly.

Autonomy changes the economics of a bad decision path

The governance problem becomes more consequential as procurement agents take on longer chains of work. A sourcing agent may search the market, retrieve supplier records, compare proposals, analyze contractual terms, invoke external services, coordinate with specialist agents and prepare an approval package before a human sees the result. Much of that activity occurs before money changes hands, but it is not economically free.

If the interpretation governing the procurement decision has already moved outside the organization's authorized position, allowing the workflow to continue creates a second problem. That organization is now consuming compute, model calls, retrieval services and agent activity on a decision path that should no longer be progressing. The earlier the organization catches that divergence, the more of the downstream compute, model calls and agent activity it never has to spend.

This matters particularly in procurement because efficiency is one of the principal reasons organizations are interested in agentic systems in the first place. Economic benefit should be measured against more than labor removed. It should also account for how effectively the architecture prevents autonomous systems from continuing to consume resources after the

© 2026 Maureen Doyle-Spare.

decision state has ceased to be one the organization would authorize.

Procurement needs a pre-execution control boundary

The durable approach is narrower: govern only the procurement decisions where the agent has real interpretive discretion, before it acts. That does not require a human approval in front of every agent action, and it does not mean relying on more logs after the transaction; routine decisions that remain within the authorized baseline can continue, while material divergence can be held or escalated.

Evidence preserved at that point can support audit, control testing and later review without turning every procurement action into an exception process. The objective is to keep autonomy inside the authority the organization intended to delegate, without slowing it down.

Agentic procurement will ultimately be judged on more than whether it can move a requisition faster or negotiate a better price. The harder institutional issue is whether an organization can allow autonomous systems to exercise meaningful procurement discretion without quietly transferring authority over what its policies mean. Approval controls still do necessary work. Before an agent spends money, selects a supplier or commits the organization to a course of action, the interpretation driving that action has to be one the organization is prepared to own. ◆

References

1 *Procurement Magazine. "Agentic AI in Procurement: Gartner, Bain, Unilever & Revolut." Procurement Magazine, August 23, 2026.* procurementmag.com/news/agentic-ai-in-procurement-gartner-bain-unilever-revolut

2 *Gartner. "Gartner Forecasts Supply Chain Management Software with Agentic AI Will Grow to $53 Billion in Spend by 2030." Press release, April 7, 2026.* gartner.com/en/newsroom/press-releases/2026-04-07-gartner-forecasts-supply-chain-management-software-with-agentic-ai-will-grow-to-53-billion-in-spend-by-2030

3 *Doyle-Spare, Maureen. Agentic AI Systems Governance: A Runtime Reference Architecture for the Reasoning Layer and the Semantic Control Plane in Regulated Financial Institutions. Zenodo, 2026.* doi.org/10.5281/zenodo.20749051

4 *Doyle-Spare, Maureen. The Semantic Deviation Index (SDI): A Runtime Measurement Standard for the Governance of Agentic AI in Financial Services. SSRN No. 6531238, April 2026.* doi.org/10.2139/ssrn.6531238

5 *Doyle-Spare, Maureen. The Economics of Agentic AI: Runtime Governance as a Distinct Determinant of Enterprise AI Cost. SSRN, July 16, 2026.* doi.org/10.2139/ssrn.7129939

Originally published in *The Agentic AI Governance Playbook* on Substack, August 23, 2026. maureendoylespare.substack.com

© 2026 Maureen Doyle-Spare.



Before an agent spends money, the interpretation driving that action has to be one the organization is prepared to own.

THE AGENTIC AI GOVERNANCE PLAYBOOK

Download the research

Papers, briefs and reference architecture at maureendoylespare.com

Visit the Playbook

WEEKLY NEWSLETTER

Follow the newsletter

The Agentic AI Governance Playbook Weekly Newsletter

Follow Newsletter

INTELLECTUAL PROPERTY: © 2026 Maureen Doyle-Spare. All rights reserved. No part of this article, including its text, figures, tables, and images, may be reproduced or distributed without the author's prior written permission.