

ENTERPRISE AI · INSTITUTIONAL MEANING

Agentic AI and *Institutional Meaning*

Why preserving organizational knowledge is no longer enough

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

Agentic AI and *Institutional Meaning* in the Enterprise

Why preserving organizational knowledge is no longer enough

MAUREEN DOYLE-SPARE

Every organization has a few people who seem to know how things really work. They understand why a process grew the way it did, when an exception deserves a second look, and why two cases that look identical should be handled differently. They carry the reasoning behind decisions made years ago, long after the projects and committees and slide decks that produced them have been forgotten. Most of that understanding was never written down anywhere, and over time it simply became part of how the institution ran.

When companies start rebuilding their operating models around agentic AI, there is an easy assumption that this kind of organizational knowledge will move over with the technology. The policies are available to the model. The procedures can be retrieved. The historical records can be searched. On paper the information is all there, ready to use. What I have found, again and again, is that having the information available is not the same as preserving what the institution actually means.

Across large transformation programs and enterprise modernization work, the technical integration is rarely the hardest part. Connecting systems, exposing APIs, and consolidating data usually turns out to be more tractable than getting different parts of an organization to agree on what an important business term really means. A customer can be approved in one system and only conditionally approved in another. A case marked complete in one workflow can still be waiting on review somewhere else. Each definition is perfectly correct inside its own context, yet the institution has never sat down and reconciled them across the enterprise. People bridge those gaps almost without thinking, because they understand the institution sitting behind the systems. An agent has none of that background to draw on.

Having the information available is not the same as preserving what the institution actually means.

Agents no longer retrieve meaning. They decide it.

As agentic systems begin reasoning across many

© 2026 Maureen Doyle-Spare.

applications at once, they end up resolving those differences on their own. They are no longer pulling back a definition someone wrote. They are deciding which working interpretation should govern the very next action they take. I call that resolved, in-the-moment reading the Operational Interpretation, and it sits at the center of my work. The interpretation an agent lands on can be completely reasonable. It can even produce the outcome most people would have expected. The trouble is that it may not be the interpretation the institution ever meant to authorize, and that is exactly where the next governance challenge begins.

For decades, enterprise governance concentrated on two moments: reviewing systems before they went live, and checking outcomes after they ran. Traditional software executed logic someone had written in advance. Machine learning widened the conversation, because statistical models needed new kinds of validation. Agentic AI changes something more basic than either of those shifts. Operational meaning now forms at runtime, before action, while the system reasons across policies, procedures, data, and context to decide what to do next.

That governance question moves with it. The challenge is no longer whether an organization can explain the output after the fact. The real question is whether it can show that the interpretation the agent acted on stayed consistent with the meaning the institution authorized in the first place. Those two things are not automatically the same and treating them as if they were, is how well-run institutions get surprised.

This reaches across every regulated industry

The pattern shows up anywhere meaning has been built up over years. A hospital depends on consistent readings of clinical guidance. A pharmaceutical manufacturer depends on shared meaning across its quality systems, validation records, and manufacturing processes. Industrial manufacturers rely on common interpretations of engineering tolerances and safety procedures. Government agencies, insurers, utilities, and critical infrastructure operators all run on operational definitions that accumulated through years of policy, regulation, and hard-won experience. The technology keeps changing, but the governance challenge underneath it stays remarkably constant.

Organizations usually describe this accumulated understanding as institutional knowledge. I have come to believe the more important asset is institutional meaning. Knowledge tells you what has been documented. Meaning determines how that knowledge gets applied at the moment an autonomous system reaches a consequential decision. That is the problem my Semantic Control Plane[1] was designed to address.

© 2026 Maureen Doyle-Spare.

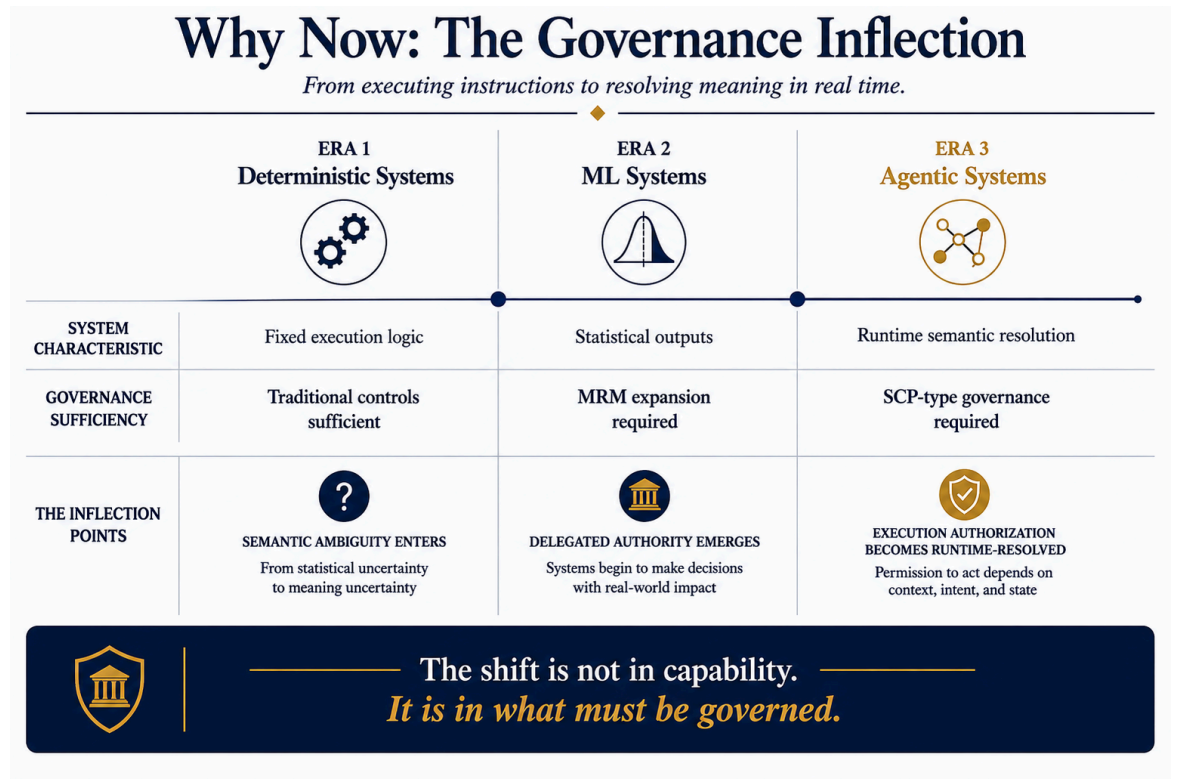


Figure 1. *Why now: the governance inflection point for agentic AI. Source: Doyle-Spare (2026)* [1]

Instead of assuming an agent will hold on to institutional intent simply because it can reach the right documents, the architecture establishes an authorized Reasoning Baseline before execution begins. Runtime interpretations get evaluated against that baseline before the system is allowed to act. The point is not to box in intelligent reasoning or strip out judgment. The point is to make sure that when an autonomous system does exercise judgment, it stays inside the operational meaning the institution has actually authorized.

My Semantic Deviation Index carries that idea further by measuring how far an agent's runtime interpretation has moved from the authorized baseline, while a Deterministic Gate enforces the response before that interpretation becomes an action. Governance stops being a review of what the agent did afterward and starts being an evaluation of the meaning it acted under, checked before the decision ever becomes operational. To my mind that shift may turn out to be one of the defining architectural changes of enterprise AI.

© 2026 Maureen Doyle-Spare.

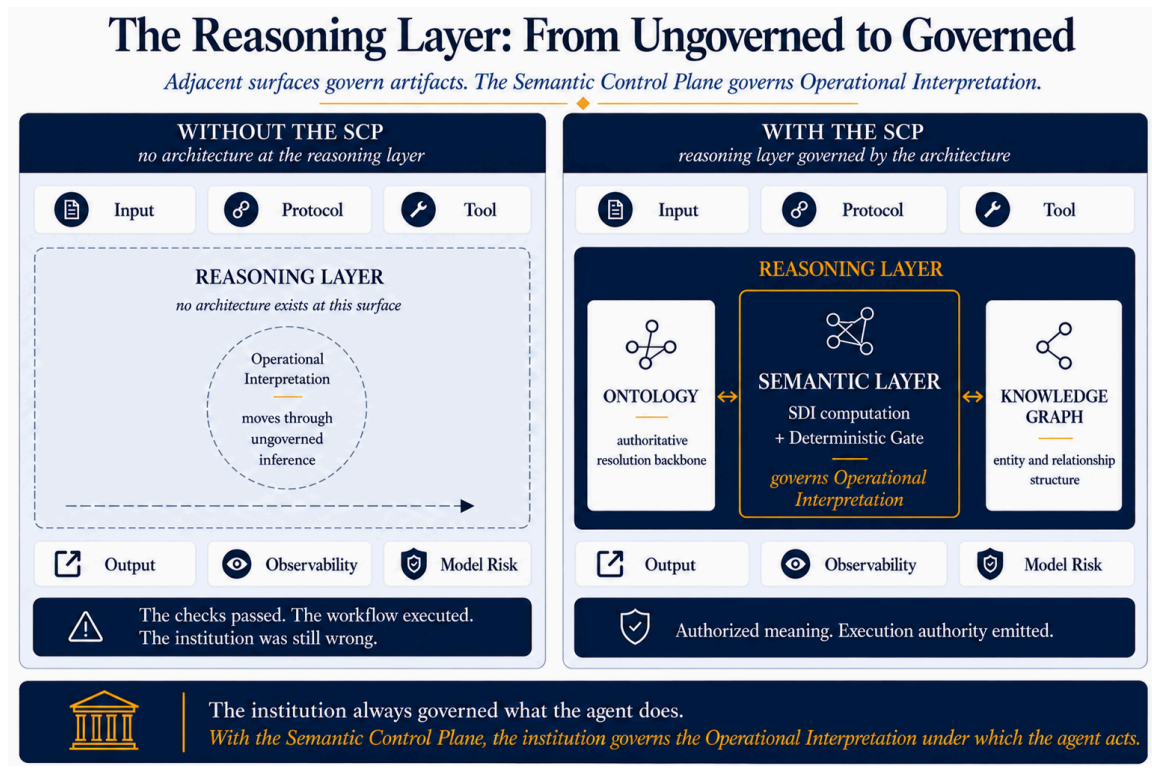


Figure 2. *The Reasoning Layer, moving from ungoverned to governed. Source: Doyle-Spare (2026)*[1]

Where this is heading

Organizations will keep investing in larger models, more capable agents, and workflows that run with less and less human involvement. Those investments will matter and they will pay off. As reasoning itself becomes part of the operating model, though, preserving institutional meaning becomes every bit as important as preserving institutional knowledge. The next generation of enterprise AI will not be judged only by how well autonomous systems reason. It will be judged by how confidently institutions can show that those systems keep reasoning inside the meaning they intended. That has quietly stopped being a technology objective, and it is becoming one of the central questions of enterprise governance. ◆

References

1 *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

Originally published in *The Agentic AI Governance Playbook* on Substack, August 2, 2026. maureendoylespare.substack.com

© 2026 Maureen Doyle-Spare.



It will be judged by how confidently institutions can show that those systems keep reasoning inside the meaning they intended.

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.