GOVERNANCE NOTE

Agentic AI and the Governance of Delegated Authority

A Governance Note on Delegated Machine Authority, Decision Rights, and Oversight Accountability

Publication Code	PD-NOTE-004
Version	1.0
Published	June 2026
Category	Cyber Risk Governance & Accountability™ (CRGA™)
Issued By	Praesidium Governance, Inc.
Canonical URL	praesidiumoversight.com/publications/governance-notes/PD-NOTE-004/
PDF Download	praesidiumoversight.com/publications/governance-notes/PD-NOTE-004_Agentic_AI_Governance_Delegated_Authority_v1_0.pdf

Document Status	Governance Note
Authority Level	Advisory
Applicability	Praesidium Governance Framework
Supersession	This document may be revised by subsequent Praesidium publications.

EXECUTIVE SUMMARY

Agentic AI does not merely automate tasks. It redistributes authority inside the enterprise.

As organizations move from AI awareness to applied agentic workflow design, the governance question changes. The issue is no longer limited to whether an AI system may generate useful outputs. The issue becomes whether software is being permitted to initiate action, recommend decisions, sequence activity, escalate issues, interact with systems, or influence enterprise outcomes.

Once AI-enabled systems can act within operational workflows, governance must define who authorized that delegation, what boundaries apply, when human intervention is required, and what evidence proves oversight occurred.

Under Cyber Risk Governance & Accountability™ (CRGA™), agentic AI is not only a technology implementation issue. It is a governance architecture issue involving delegated machine authority, decision rights, escalation discipline, accountability boundaries, and defensible oversight evidence.

The Development

Enterprise AI adoption is moving quickly from general awareness and policy discussion into applied workflow redesign.

Organizations are no longer asking only whether employees may use AI tools for productivity, research, drafting, analysis, or summarization. They are increasingly evaluating AI-enabled workflows that may recommend action, trigger process steps, interact with enterprise systems, support customer-facing activity, analyze operational data, prioritize risk, or coordinate activity across business functions.

This shift is especially important in the rise of agentic AI. Unlike static automation or narrow decision-support tools, agentic systems may operate across multi-step workflows. They may interpret goals, select actions, interact with other systems, route information, generate recommendations, escalate exceptions, or execute tasks with varying levels of human supervision.

That movement creates a new governance condition.

When software is used to assist a human, the governance question is often one of use, accuracy, privacy, and acceptable reliance. When software is permitted to act within an enterprise workflow, the governance question becomes one of delegated authority.

The distinction matters.

- A chatbot may create output.
- An agentic workflow may alter action.
- A delegated system may affect authority.

The Governance Issue

Agentic workflows create governance exposure because they may alter decision rights, escalation pathways, system access, and human accountability.

In traditional governance models, authority is usually attached to people, roles, committees, policies, approval thresholds, and reporting structures. Those structures define who may make decisions, who must be informed, who may approve exceptions, and who is accountable when a material issue arises.

Agentic AI complicates that structure because software may begin operating inside the space where authority is exercised.

- A system may recommend which risk receives priority.
- A system may escalate one issue while suppressing another.
- A system may initiate a workflow that creates operational, customer, financial, legal, cyber, or reputational consequences.
- A system may interact with sensitive data, privileged systems, third-party platforms, or customer-facing processes.
- A system may make human review appear present when, in practice, human review is delayed, superficial, or absent.

The governance problem is not simply that AI may be inaccurate. The deeper issue is that AI may be permitted to act without a clearly governed delegation model.

For boards and executives, the relevant questions are not only technical questions. They are governance questions:

- Who authorized the workflow?
- What authority was delegated?
- What remains reserved to humans?
- What escalation thresholds apply?
- Who owns exceptions, failures, and downstream harm?
- What evidence proves that oversight occurred before, during, and after deployment?

Without those answers, agentic AI can create a gap between operational activity and accountable authority.

That gap is where governance exposure forms.

Why Existing AI Policies Are Not Enough

Many organizations have begun developing AI policies, acceptable-use standards, employee guidance, vendor review procedures, and model risk controls. These are useful, but they are not sufficient for agentic AI governance.

- A policy may state what employees are allowed to do.
- It may define prohibited uses.
- It may address data handling, privacy, confidentiality, bias, intellectual property, or human review.
- It may require approval before deploying certain AI tools.

But an AI policy often does not define the full governance architecture required when AI is embedded into enterprise workflows.

Specifically, many AI policies do not clearly answer:

- What authority may be delegated to an AI-enabled workflow?

- Who has authority to approve that delegation?
- Which decisions, actions, or exceptions must remain human-controlled?
- What systems, data environments, or privileged access pathways may the AI interact with?
- What escalation thresholds apply when the AI detects, recommends, or initiates action?
- What evidence must be retained to prove executive oversight?
- What conditions require suspension, rollback, review, or board notification?

This is why AI governance cannot stop at policy issuance.

A policy may define acceptable use. Governance architecture defines accountable authority.

Agentic AI requires the second.

CRGA™ Interpretation

Under CRGA™, agentic AI belongs within AI & Algorithmic Risk Governance, but it also intersects with identity, access, cyber, data, third-party reliance, and enterprise accountability.

This is because agentic workflows rarely remain confined to a single technology domain. They may depend on data access, identity permissions, privileged system interaction, third-party platforms, automated decision pathways, workflow orchestration, and human-agent operating structures.

As a result, agentic AI should not be governed only as an AI ethics issue, a compliance issue, a security issue, or an IT implementation issue. It should be governed as a technology-enabled enterprise risk condition where delegated machine authority may materially affect the organization.

CRGA™ treats this as a governance architecture question. The relevant oversight layer must define:

- Decision rights — who may approve agentic workflows and delegated machine authority.
- Accountability boundaries — who remains responsible for system behavior, workflow outcomes, exceptions, and failures.
- Escalation thresholds — when issues must move from operational handling to executive, risk, legal, compliance, audit, or board attention.
- Authority limits — what the system may recommend, initiate, execute, or escalate without additional human approval.
- Evidence requirements — what records demonstrate that governance occurred before deployment and during operation.
- Review conditions — what events require reassessment, suspension, rollback, or board-level reporting.

The issue is not whether agentic AI can create operational value. It can.

The issue is whether the organization can prove that delegated machine authority was governed before it was exercised.

Board-Level Questions Before Agentic AI Goes Live

Before an organization deploys agentic AI into a material workflow, boards and executive leadership should be able to answer seven governance questions.

1. What authority is being delegated to the AI system?

The organization should define whether the AI system is merely generating information, recommending action, initiating workflow steps, interacting with enterprise systems, escalating issues, or executing tasks. Different levels of autonomy create different governance obligations.

2. Who approved that delegation?

Delegated machine authority should not emerge informally through tool adoption, vendor configuration, departmental experimentation, or workflow convenience. The organization should identify the accountable approval body, executive owner, or governance forum responsible for authorizing the delegation.

3. What decisions or actions remain reserved to humans?

Human review should not be treated as a vague safeguard. The organization should define which decisions require human approval, which exceptions require human intervention, and which actions may not be performed by an AI-enabled system under any circumstance.

4. What escalation thresholds apply?

Agentic workflows should have defined escalation thresholds for risk, uncertainty, anomalies, exceptions, regulatory exposure, customer impact, system access, security events, financial consequences, and reputational sensitivity. The system should not only act. It should know when not to act.

5. What evidence will prove oversight occurred?

Governance must be evidenced. The organization should retain records showing approval, risk evaluation, authority boundaries, access limitations, testing, monitoring, exception handling, and periodic review. Absent evidence, the organization may have activity but not defensible oversight.

6. Who owns failures, exceptions, and downstream harm?

Accountability cannot be delegated to software. The organization should identify the executive, business function, control owner, risk owner, or governance body accountable for failures, exceptions, unintended consequences, and downstream harm.

7. What conditions require suspension, rollback, or board notification?

Agentic AI governance should include predefined conditions that require the workflow to be paused, modified, escalated, or reported. These conditions should be established before deployment, not improvised after failure.

Praesidium Position

Agentic AI becomes governable only when delegated machine authority is bounded by decision rights, escalation discipline, accountability structures, and defensible oversight evidence.

The enterprise risk is not simply that AI systems may produce flawed outputs. The deeper risk is that organizations may permit software to operate inside material workflows without first defining the governance conditions under which

authority may be delegated.

For boards and executives, the central issue is not whether the organization is adopting AI quickly enough. The central issue is whether the organization is adopting AI in a way that remains governable when action is taken, authority is exercised, and consequences arise.

- AI policies are necessary.
- Technical controls are necessary.
- Vendor reviews are necessary.
- Security and privacy safeguards are necessary.

But they are not substitutes for governance architecture.

As agentic AI becomes embedded into enterprise workflows, organizations will need to demonstrate more than innovation. They will need to demonstrate that authority was assigned, bounded, escalated, reviewed, and evidenced.

That is the governance obligation.

Under CRGA™, agentic AI is not merely an automation trend. It is a test of whether enterprise governance can keep pace with delegated machine authority.

CLOSING OBSERVATION

Agentic AI changes the oversight question.

The issue is no longer only what AI can produce. The issue is what AI is allowed to do.

Where software can initiate, recommend, sequence, escalate, or execute action, governance must precede deployment. Decision rights must be defined. Human authority must be preserved where required. Escalation thresholds must be clear. Accountability must remain assigned. Evidence must be retained.

Agentic AI does not remove the need for human governance.

It makes governance more urgent.