Risk	Contractual Mitigation	Operational Mitigation
Unpredictable Decision-Making: Agentic AI systems can make choices that deviate from expected parameters.	• Service descriptions and security mechanisms must define acceptable operating boundaries. • Include provisions that limit actions outside predefined parameters.	• Create organizational policies that define acceptable boundaries. • Limit the APIs that enable interaction with critical systems. • Implement real-time monitoring systems and escalation mechanisms to identify and address deviations quickly.
Bridging the Techno-Responsibility Gap: If an AI agent causes harm, liability may need to be allocated according to the delivery model.	• Ensure that the vendor remains responsible for: – Inherent design flaws – Inadequate safety controls – Knowable risks not disclosed during procurement – Failures in core system architecture, or failure to implement updates or patches for known issues – Violations of promised performance parameters • Define a clear and actionable acceptable use policy with which the customer remains exclusively responsible for compliance.	• Establish governance frameworks and conduct regular risk assessments to address shared and emergent risks. • Ensure cyber insurance covers risks associated with the use of agentic AI.
Greater Need for Human Oversight: Agentic AI often requires ongoing monitoring to ensure compliance, avoid boundary drift, and prevent unauthorized actions.	• Require vendors to build in human-in-the-loop functionality. • Require mechanisms for prompt human intervention or override in critical scenarios.	• Designate responsible human operators to monitor, intervene, and ensure compliance with ethical and legal standards. • Adequately, and periodically train employees responsible for monitoring the tool.
Failure of the Intended Purpose (Interoperability): Agentic systems may interface with third-party APIs or systems autonomously.	Mandate robust interoperability standards and vendor warranties for integration with third-party systems.	Test integrations periodically.
Termination Protocols: Given agentic AI's autonomy, shutting it down requires more than pulling the plug. Systems must be de-integrated. However, the de-integration process cannot inhibit prompt suspension of agentic actions if issues arise.	• Prescribe technical, operational, and legal procedures for deactivation without disruption. • Address data ownership and retrieval upon termination or suspension. • Require the implementation of mechanisms to pause, halt, or override decisions when specific thresholds are breached.	Prepare a detailed termination and suspension playbook, including safe decommissioning procedures and contingency planning.
Lack of Transparency: Given the black box nature of most AI tools, understanding how or why the agent performs in a certain way may be indiscernible.	Require explainability and auditability standards to clarify decision-making processes	• Perform regular assessment of decision-making patterns. • Prepare teams to work alongside autonomous systems, emphasizing collaboration and oversight.