

# LEAH



## Agentic AI Use Cases for Legal Teams

From AI Assistance to Autonomous Legal Execution

---

# Executive Summary

Legal departments are being asked to do more—with less. Contract volumes are rising. Business teams expect faster turnaround. Regulatory scrutiny is intensifying. Global instability is increasing. Boards expect sharper risk visibility. And headcount is not scaling at the same rate as demand.

Traditional AI tools have delivered incremental gains. Document review, clause extraction, and research automation have improved productivity at the individual level. But most of these systems remain assistive; they summarize, extract, and suggest. They do not execute or orchestrate.

## A new paradigm is emerging: Agentic AI.

Agentic AI moves beyond single-task automation to orchestrating multi-step workflows. It reasons across systems. It enforces playbooks. It routes approvals. It maintains audit trails. It adapts to context. It operates inside defined guardrails while keeping humans firmly in control.

This is not an incremental innovation. It is a structural shift in how legal work gets executed:

- Contracts that move from inbox to redline automatically within policy guardrails
- Compliance reviews that trigger themselves on cadence and regulatory change
- Obligations that become managed workflows rather than forgotten clauses
- Governance processes that are continuously audit-ready
- Diligence exercises that surface risk systematically and defensibly

# 33%

of enterprise software apps will integrate agentic AI by 2028 (Gartner).



Importantly, Agentic AI does not replace legal judgment. It elevates it. It removes the mechanical burden of execution so legal professionals can focus on strategy, negotiation, and risk calibration—the work that requires human expertise.

For General Counsel and Legal Operations leaders, the question is no longer whether to adopt AI. It is whether your legal function will remain reactive and workflow-bound or evolve into a governed, orchestrated intelligence engine that scales without proportionally scaling headcount.

This report explores what agentic AI actually means in the context of legal operations—beyond the hype, beyond copilots, and beyond surface-level automation. It explores practical, high-impact Agentic AI use cases. For each use case, it outlines how legal teams can start with targeted AI assistance, build confidence through structured automation, and progressively expand into fully orchestrated workflows, scaling capability and governance maturity over time.

The organizations that adopt Agentic AI thoughtfully will not just work faster. They will operate differently: with greater consistency, transparency, and strategic influence across the enterprise.

# Table of Contents

---

Evolution of Enterprise AI	5
Stages of AI Maturity	7
Procurement Use Cases	9
Contract Management	10
Obligation Management	12
Regulatory and Compliance Oversight	14
Corporate Governance	16
Enterprise Risk and Due Diligence	18
The Path Forward	20

---

# Evolution of Enterprise AI

Enterprise AI isn't following a linear path of gradual improvement. It's undergoing three distinct phase shifts, each fundamentally changing what's possible in how businesses operate. Each phase unlocks different operational possibilities and requires different organizational capabilities to capture value.

---

## Phase 1

Generative AI  
(User-Led)

**This first wave delivered reactive, single-step capabilities with no workflow memory or external integration.** AI operated as an on-demand content generator—fast and impressive for discrete tasks like drafting documents or summarizing documents, but fundamentally limited. Users had to initiate every action, provide all context, and manually connect outputs to broader workflows. The impact ceiling was individual productivity, not organizational transformation.

---

## Phase 2

Agents  
(Goal-Led)

**The second phase introduced agents capable of proactive, multi-step execution.** These systems could maintain task-level memory, use tools and APIs to interact with other platforms, and complete workflows autonomously once given a goal. A procurement agent, for example, could gather supplier data, draft comparison analyses, and update tracking systems without constant human prompting. However, these agents still operated in functional silos, handling specific workflows rather than orchestrating cross-functional processes.

---

## Phase 3

Agentic AI  
(Business Outcome-Led)

**This emerging phase represents a fundamental shift in how enterprises operate.** Agentic AI systems don't just complete tasks or workflows, they manage business outcomes across multiple agents, functions, and systems. These autonomous platforms can coordinate procurement, legal review, financial reconciliation, and stakeholder approvals simultaneously, making contextual decisions while maintaining compliance guardrails. They engage in long-horizon planning, adapt strategies based on changing conditions, and continuously optimize for business results rather than task completion.

	Gen AI	AI Agents	Agentic AI
<b>Autonomy &amp; Decision-Making</b>	<ul style="list-style-type: none"> <li>• Requires human prompts</li> <li>• Suggests, does not act</li> <li>• Flags issues for manual handling</li> </ul>	<ul style="list-style-type: none"> <li>• Execute defined tasks</li> <li>• Make task-level decisions</li> <li>• Resolve simple exceptions</li> </ul>	<ul style="list-style-type: none"> <li>• Initiates workflows autonomously</li> <li>• Makes cross-functional decisions</li> <li>• Resolves most operational exceptions</li> <li>• Operates within governance guardrails</li> </ul>
<b>Contextual Understanding</b>	<ul style="list-style-type: none"> <li>• No memory beyond session</li> <li>• Processes information in isolation</li> <li>• Limited awareness of business context</li> </ul>	<ul style="list-style-type: none"> <li>• Maintain task-level memory</li> <li>• Understand workflow context</li> <li>• Apply structured logic</li> </ul>	<ul style="list-style-type: none"> <li>• Persistent cross-workflow memory</li> <li>• Enterprise-wide contextual awareness</li> <li>• Aligns to policy, risk thresholds, KPIs</li> </ul>
<b>Integration &amp; Orchestration</b>	<ul style="list-style-type: none"> <li>• No system execution</li> <li>• Manual transfer between tools</li> <li>• Standalone outputs</li> </ul>	<ul style="list-style-type: none"> <li>• API-enabled execution</li> <li>• Multi-step predefined workflows</li> <li>• Connect specific systems</li> </ul>	<ul style="list-style-type: none"> <li>• Orchestrates multiple systems dynamically</li> <li>• Coordinates across legal, procurement, finance, sales</li> <li>• Enables multi-agent collaboration</li> </ul>
<b>Learning &amp; Adaptability</b>	<ul style="list-style-type: none"> <li>• Remains largely static</li> <li>• Requires manual prompt refinement</li> </ul>	<ul style="list-style-type: none"> <li>• Improve through workflow tuning</li> <li>• Performance improves with task over time</li> <li>• Adjust within defined rules systems</li> </ul>	<ul style="list-style-type: none"> <li>• Continuously improves through operational feedback loops</li> <li>• Optimizes across workflows and outcomes</li> <li>• Dynamically adapts within governance guardrails</li> <li>• Builds operational intelligence over time</li> </ul>

# Stages of AI Maturity

## Stage 1

### Prompt-Driven Assistant

(AI assists individuals)

**At the foundation level, AI is used primarily as an assistant.** It generates drafts, summarizes documents, answers questions, and supports brainstorming. The interaction is prompt-based and session-based, meaning AI acts only when asked and does not retain workflow context over time.

Humans remain fully responsible for decisions, execution, and follow-through. The impact is typically limited to individual productivity gains—faster writing, quicker research, or improved idea generation—without structural changes to how the business operates.

## Stage 2

### Task Executor

(AI completes defined work)

**AI shifts from content generation to action.** Agents can execute structured, multi-step tasks such as gathering information from systems, drafting documents based on templates, updating CRM or ERP records, routing requests for approval, or compiling reports.

Humans still define the objective and review the final output, but much of the repetitive execution is automated. This stage reduces manual workload, increases consistency, and shortens turnaround times. However, AI remains reactive and limited to specific tasks rather than owning broader processes.

## Stage 3

### Workflow Owner

(AI runs processes end-to-end)

**AI moves beyond completing tasks and begins managing full workflows.** Agents monitor deadlines, pull data from connected systems, route approvals, flag risks, and maintain context across long-running processes such as contract reviews, purchase approvals, or financial reconciliations.

Humans oversee performance and step in for exceptions, escalations, or strategic judgment calls. Instead of manually pushing work forward, teams supervise and refine processes. The result is meaningful cycle-time reduction, fewer bottlenecks, improved visibility across departments, and more predictable operational performance.