



Agentic AI – the transition to Agent Orchestration



Presentation by: **Jeffrey Seymour**
Date: September, 2025

Defining the modern AI Agent


AI agents are software entities powered by large language (or multimodal / reasoning) models that **once given a goal** they...



- 01 Can plan the steps
- 02 Invoke external tools and data
- 03 Execute end-to-end workflows
- 04 Iteratively improve from feedback

Their **autonomy**, **intentionality**, and **adaptive learning** push them well beyond deterministic automation—turning static processes into self-optimizing systems


From goal to output: How an agent runs



Orchestration is the process of managing a multi-step workflow – deciding what needs to happen, when, and by which agent or tool.

Single Agent Orchestration


One agent performs multiple steps by itself – looping through logic, making decision, and using tools



Example: A **support agent** that reads a ticket, pulls account data, suggests a solution, and follows up—all within one loop

Multi-Agent Orchestration

A **manager agent** acts like a project manager, delegating tasks to other specialized agents and compiling the outcome.



Example: A **Manager Agent** receives a support ticket and delegates it to a **Ticket Analysis Agent**, which compiles relevant customer context

- If the issue is routine and the customer is not flagged as high risk, it routes the case to the **Standard Resolution Agent**
- If the customer shows signs of dissatisfaction the Manager Agent routes the case to a **Churn Avoidance Agent** for a more tailored and proactive resolution


Future of Claims Adjudication




> **90% time savings**

AI-driven claims adjudication automates validation, approvals, and payment processing. It drafts EOBs and appeal materials, ensuring **faster resolutions and lower cost of labor**, and an **improved experience**.

From today's manual process (9 steps driven by a human)...



... to an agent-driven future (2 steps driven by a human)



Risks of Agent Orchestration

Transitioning from single agent use to orchestrated multi-agent systems introduces new risks:

Unintended Data Access

- Multi-agent workflows increase the risk of unauthorized data exposure and leakage.
- Complex permissions and inter-agent communication can bypass traditional controls.

Unexpected Behaviors

- Autonomous agents may act in conflicting or unpredictable ways.
- Emergent dynamics can amplify errors or create feedback loops.

Lack of Transparency

- Decision chains become opaque, making it difficult to audit or explain outcomes.
- Reduced human oversight increases the risk of undetected errors or bias.

Mitigating Orchestration Risks

Key strategies to address orchestration risks:

- Implement strong governance and guardrails (agent registries, continuous monitoring).
- Prioritize transparency and observability in agent interactions.
- Maintain human oversight for critical decisions and edge cases.
- Regularly review and update access controls and audit trails.

Thank you