

# AI Agents: What They Are, What They're Not, and How to Deploy Them Without Burning Cash Part 2

What AI Agents Really Are: An Executive Briefing

## From Promise to Practical Value



Last week I shared part 1 “The Three Misconceptions Costing You Money”. This article continues the exploration of this topic.

Strip away the jargon and an AI agent does three things:

1. **Perceives**, it monitors its environment through data feeds, APIs, databases, sensor inputs, market signals, or any other digital source you connect it to.
2. **Processes**, it reasons about what it has observed, makes decisions, and plans a course of action. This is where the LLM backbone earns its keep.
3. **Acts**, it executes actions in your systems. Triggering workflows, updating databases, sending notifications, generating reports, or calling other agents.

That perceive-process-act loop is what makes an agent an agent. It is not magic. It is a software system with an AI backbone and defined autonomy.

### Three Types, Three Business Applications

Not all agents are built the same. The type you need depends on the problem you are solving.

**Reactive agents** respond directly to inputs with no memory or planning. They are fast and lightweight. Use them for status monitoring, simple alerts, and threshold-based triggers. Think of a system that watches your server infrastructure and pages your team when CPU usage crosses 90%. **Fast, reliable, narrow.**

**Deliberative agents** use reasoning, planning, and long-term memory. They are slower but far more capable. Use them for research tasks, strategic analysis, and complex decision support. Think of a system that analyses three months of market data, identifies emerging trends, and produces a competitive intelligence briefing. **Thorough, nuanced, resource-intensive.**

**Hybrid agents** combine reactive speed with deliberative depth. They handle the bulk of production workloads. They react quickly to straightforward situations and escalate to deeper reasoning when the situation demands it. Most of what you will deploy in practice will be hybrid.

## The Supervisor Pattern: Your Digital VP

The most effective multi-agent architecture for business use follows what engineers call the Supervisor pattern. It works exactly like a well-run department:

A user or system submits a request. A Supervisor Agent receives it, assesses what needs to happen, and delegates to the right specialist agents. Those agents do their work and report back. The Supervisor assembles the final output.

Picture your compliance department. A regulatory change is published. The Monitor Agent detects it. The Research Agent pulls the full text and related guidance. The Analysis Agent maps it against your current policies and identifies gaps. The Execution Agent drafts updated procedures. The Supervisor Agent reviews the whole package and routes it for human approval.

That is not science fiction. That is available today. I have built a few and continuing to do so.

**Bottom line:** An AI agent is a software system that perceives, reasons, and acts. Multi-agent systems are teams of these, with clear roles and a supervisor. Deploy the type that matches your problem.

## The Organisational Earthquake Most Executives Are Not Prepared For

This is the section most AI articles skip. And it is arguably the most important one.

Previous waves of technology, automation, outsourcing, cloud computing, attacked tasks. They made specific activities faster or cheaper. But humans retained their advantage because humans are goal-seeking. We plan, prioritise, sequence, and adapt. Technology could handle individual tasks, but the coordination of tasks towards objectives remained a human capability.

AI agents break that pattern.

Agents are goal-seeking technologies. They do not just execute a task you assign them. They take a goal, decompose it into tasks, plan the sequence, execute against it, and adapt when things change. This is the first time technology can substitute not just tasks but actual goals.

And that changes everything about how organisations work.

## Roles Are Bundles of Goals

Consider any role in your organisation. It is not a single activity. It is a bundle of goals. Your head of market research has multiple objectives, competitive intelligence, customer insight, trend analysis, pricing strategy. Each of those goals contains dozens of underlying tasks.

When previous technologies automated tasks within those goals, the role stayed intact. The person still owned the goal. They just had better tools.

When an AI agent takes over an entire goal, say, competitive intelligence, that goal is unbundled from the role. The scope of the role shrinks. The person retains their other goals, but their portfolio is smaller.

Scale that across an organisation and you get a cascading effect. Roles with reduced scope get merged. Teams that depended on a specific role's contribution to achieve their objectives can now use an agent instead. The person keeps their job, but their organisational influence, their cross-team visibility, their negotiating power, all of that erodes.

As one analyst put it: "**AI will not take your job, but it will take your ability to charge a premium for it.**"

### **This Is Not About Layoffs. It Is About Restructuring.**

The executives who frame AI agents purely as a headcount reduction tool are missing the point. The real impact is structural.

Agents will force a fundamental rebundling of roles and teams. Some roles will absorb the remaining goals of roles that have been hollowed out. Some teams will shrink as agents replace specific contributors. New roles will emerge, agent architects, workflow designers, human-AI coordination managers.

We have seen this dynamic before with outsourcing and offshoring. Goals that were previously performed in-house could be performed externally. But those were always human resources. The rate at which they learned and adapted was limited. AI agents learn faster, embed deeper, and cost less at scale.

The organisations that thrive will not be the ones that deploy the most agents. They will be the ones that redesign their structures around a hybrid workforce of humans and agents, with clear boundaries for what each does best.

**Bottom line:** AI agents do not just automate work. They reorganise it. If you are not planning for structural change, you are planning to be surprised by it.

#ai-agents #business-strategy #c-suite #agentic-ai #enterprise-ai

## Disclaimer and Disclosure

**Third-party Content and AI Assistance:** This article references tools and software that are publicly available and proprietary to their respective creators. The author does not claim ownership or affiliation with these third-party products. This article was written by the author with assistance from Generative AI Language Models.

**Transparency Notice:** While every effort has been made to ensure accuracy, readers should verify information independently and consult official sources or documentation for the mentioned tools and software. The use of AI in the writing process is disclosed in the interest of transparency, but all opinions and analyses are the author's own unless otherwise stated.

