AI BASICS // AGENT DESIGN

Designing agentic AI systems that actually deliver

The real opportunity isn't building agents. It's how you think about them, and that starts with how you architect the system.

> **Travis's short version:** don't build one mega agent. Design a system where each part has a clear role, and work through four stages: intent, scouting, synthesis and output. The goal is practical, actionable insight, not something that's simply novel.

A simple mental model

Before we get into the details, this is the mental model we use when designing agent systems. It keeps things grounded: each agent has a job, and the whole system flows towards one outcome. To bring it to life, we walk through a scenario you could design today: an agentic workflow that surfaces the latest AI trends across industries.

Stage	The question it answers	What the agent does
1. Intent	What are we actually trying to do? What is the user trying to achieve?	Doesn't take the prompt at face value. Asks whether it needs more context and loops back with clarifying questions.
2. Scouting	Where do we get the best information from?	Each scouting agent owns one domain, such as AI in healthcare, fintech or retail, and pulls from a curated, high-quality list of sources.
3. Synthesis	What does it all mean?	Reads across everything, looks for patterns, highlights weak signals, then checks its own work and keeps refining.
4. Output	Who needs this, and how should it be delivered?	Packages the insights as a quick summary, a lightweight report or a visual dashboard, whichever suits the context.

Stage one: nail the intent

This is where most systems go wrong. A typical agent takes "find AI trends" at face value. Great, but what kind of trends, from where, how often and for whom? Most people can't fully articulate what they want on the first go, and that's completely okay. We do the same ourselves.

That's why we start with an **intent agent**. It asks itself, "Do I need more context to get this right?" If yes, it asks things like: are you focused on global trends or a specific region, and do you want updates daily, weekly or on demand? Without nailing intent first, the whole system is built on guesswork, and that's where most projects fall apart.

Stage two: cast a tight net

Scouting agents each get their own focus and a curated list of resources. Why? Because if you let agents crawl the whole web, you end up summarising SEO sludge. That's not insight, it's internet noise. This is where subject matter experts play a key role: not by writing the prompts, but by helping to find which sources actually matter, such as newsletters, academic databases and industry portals. **Less junk, better signals, more meaningful insights.**

Stage three: think, check, repeat

Synthesis is what makes the system more than a scraper. Without it, even the best scattered data just sits there, disconnected. The synthesis agent works in a loop:

1. **Read across everything.** It looks for patterns and highlights weak signals, rather than just summarising headlines.
2. **Reflect on the result.** It asks itself whether what it produced is actually relevant. If not, it adjusts, reframes the summary and pulls in different context.
3. **Check for enough insight.** If it doesn't have enough, it keeps refining and digging deeper, looping until it's confident the output is useful.

It's not just collecting. It's thinking, adapting and getting sharper each time around.

Stage four: deliver it

Even the sharpest insights mean nothing if they don't reach the right person in the right format. The output agent turns refined insights into something usable, and stays flexible about the format. The system should highlight which insights will actually be consumed, rather than dumping everything into a doc. Great thinking that never gets delivered is just shelfware.

The bottom line

You don't need to be a developer to start thinking like this. All it takes is some structure, the right tools and a mindset that is less about "what can this agent do?" and more about **"what does the system need to deliver?"** Building intelligent agent systems isn't about complicated code or a huge budget. It's about clear design and purposeful roles, with no-code platforms having unlocked a world of possibilities.

Want to talk this through for your business?

Book a free 15-minute call. We'll look at your situation and give you a straight answer, with no jargon and no sales pitch.

Let's Chat

Opens the request form on our website. Free, no obligation.