

Mission Control: A Field Guide to Running Your AI Agents

**Everyone is building AI agents.
Who is running the mission?**

A point of view from Axon Ivy, part of the Ricoh group.

Following the session at the Gartner Application Innovation & Business Solutions Summit, London 2026.



Mission Control was never about the rocket. It was about the room.

Apollo flew because a team of people could see everything the spacecraft was doing, decide in real time, and step in the moment something drifted. The vehicle had autonomy. The mission had control.

Most companies today have the opposite. They have launched a sky full of AI agents and built almost none of the room. Agents draft, decide, route, reconcile and act across systems, and in a growing number of organisations no single person can say what all of them are doing right now.

That is not a technology problem. It is a control problem, and it is the one that decides whether AI becomes leverage or liability.

Five questions Mission Control has to answer before you can say your AI is under control.

Visibility

You cannot run what you cannot see.

Agents get deployed team by team and tool by tool. Each one is reasonable on its own. Together they become a fog. When something goes wrong, the first hour is spent working out which agent even touched the process.

What good looks like: Every agent, action and decision visible in one place and in business terms, not buried in logs.

Authority

Humans belong in the flow, not at the end of it.

The common design bolts a human approval onto the very end. By then the context is gone and the answer is almost always yes.

What good looks like: People positioned where the decision actually matters, inside the process, with the authority and information to redirect it.

Coordination

Agents that work together, not merely in parallel.

An agent that is brilliant on its own still fails if it cannot hand off cleanly to the next agent, to the core system, and to the person who owns the outcome.

What good looks like: Orchestration across the whole chain, so agents, existing systems and people run one process.

Accountability

A record of who did what, and why.

When an autonomous decision turns out wrong, “the AI did it” is not an answer a regulator, an auditor or a customer will accept. Most agent stacks cannot reconstruct the reasoning after the fact.

What good looks like: Every step traceable end to end, so any outcome can be explained, defended and improved.

Control

The ability to intervene, in flight.

Autonomy without a stop button is not confidence, it is exposure. If you cannot pause, reroute or shut down an agent while it is working, you do not run it. It runs you.

What good looks like: Real-time intervention treated as a core feature.

The room itself

Put those five together and you have a command center for autonomous business. That is the layer we build at Axon Ivy.

It does not replace the systems and agents you already run. It connects and coordinates them, keeps people in the flow, and keeps the whole process visible, accountable and under control.



Is your AI actually under control?

Six quick checks. Run them against your own business.

- ☐ Can you name every AI agent running in your business right now?
- ☐ Can you see, in one view, what all of them are doing today?
- ☐ When an agent makes a call, is a person positioned to change it before it lands, not after?
- ☐ If an outcome is challenged, can you reconstruct exactly how it was reached?
- ☐ Can you pause or reroute any agent while it is in flight?
- ☐ Does one process run across your agents, systems and people, or several disconnected ones?

Hesitating on more than one? The agents are flying. The mission is not yet under control.

The advantage will not go to whoever deploys the most agents. It will go to whoever can run them.

Everyone is building AI agents. Mission Control is the part almost no one has built yet. It is the room worth building, and it is where the next few years will be won.

Talk to our team, or see Mission Control in action.

www.axonivy.com