ESSAY · ENTERPRISE AUTONOMY

The Hard Part Was Never the AI.

Agents, robots, and the operating model most companies don't have yet.

By Dr Ramzi Ben Ouaghrem

July 2026 · 6 min read

THE AUTONOMY OPERATING MODEL — FIVE QUESTIONS

01 What work should machines perform?
02 How much autonomy should each machine hold?
03 Who governs machine authority?
04 How do we plan across human, digital and physical capacity?
05 What are we actually measuring?

Two sentences, one word apart.

> "Tell me which invoices are overdue."
>
> **"Resolve our overdue invoices."**

The first needs a system that can read your ledger. The second needs one that can log into your finance stack, decide who to chase, offer terms, update records, and in some setups move money. It also needs a named human who answers for it when it gets one wrong.

Change one verb and you have changed what the technology is, and who is responsible for it.

Most AI programmes stall at exactly this step, from answering to acting. Not because the models are not good enough. Because nobody has decided what the machine is allowed to do.

It is one shift, not two

We treat AI agents and robots as two separate agendas. Different budgets, different committees, usually different parts of the org chart. One works inside the ERP. The other works on a line.

They are the same shift: machines that no longer just process work, but perform it.

Which turns the strategic question into something less comfortable than *where can we use AI*.

> **Which decisions and actions are we willing to hand to a machine, and on whose authority?**

That is not a technology question. It is a delegation-of-authority question, and most boards have not looked at that document with machines in mind.

The constraint is not intelligence

Two findings convinced me of this.

BMW ran a humanoid robot at its Spartanburg plant for ten months. It handled more than 90,000 components across roughly 1,250 operating hours and contributed to over 30,000 vehicles. Then BMW retired it. The next pilot is in Leipzig, using a different company's machine.

What BMW kept was not hardware. It was the operating data, the integration lessons, and a new internal Center of Competence for Physical AI. BMW is also specific about what made the pilot possible at all: a unified data model across the

production system, for the robot to plug into. The humanoid was the visible part. It was not the hard part.

Microsoft's 2026 Work Trend Index gets to the same place from the software side. It tested 29 factors against measured AI impact. Two-thirds of that impact traced back to organisational conditions: culture, manager support, whether governance exists at all. The remaining third came down to individual mindset.

> **The bottleneck is not intelligence. It is whether the organisation is ready to receive it.**

An agent with write access is a worker

The moment a system acts rather than answers, questions appear that nobody on the AI team is placed to answer.

- Who authorised this?
- Whose credentials is it using?
- How much can it spend?
- What can it change?
- Can it hand work to another agent, and can you reconstruct that chain afterwards?
- Who is accountable when it is wrong?
- **And, in some cases : how do you stop it?**

These are identity and access questions, which makes them enterprise architecture rather than model safety. NIST published a concept paper in February 2026 naming agent identity and authorisation as a foundational gap. One problem it flags as unsolved is worth carrying around: single-hop delegation works today, but nobody has a settled answer for what happens when one agent authorises another.

So my position is straightforward. An agent with write access to production systems is a worker holding delegated authority. Give it an identity, a scope, a spending limit, a supervisor, an audit trail, and a named human who answers for it. We would not take on a contractor under looser terms.

Robots need all of that, and one thing more. A physical action often cannot be undone. A bad slide deck gets rewritten; a robot that puts a component through a windscreen has made a permanent contribution to the day. Which is why a robot's limits have to be deterministic and physical, not implemented in a prompt.

Autonomy is a ladder, not a switch

Every rung moves authority from a person to a machine.

THE AUTONOMY LADDER			AUTHORITY TRANSFERRED →
1	**Assist**	Provides information or drafts content. The human does the work.	
2	**Recommend**	Proposes a decision. A human decides and executes.	
3	**Act with approval**	Prepares and executes, but only after human authorisation.	
4	**Act within bounds**	Operates independently inside predefined limits.	
5	**Orchestrate**	Plans and coordinates other agents, systems and machines toward an outcome.	

Target the highest rung you can credibly govern, not the highest rung available.

The mistake will be treating level five as the destination for everything. The right target is the highest rung you can actually govern, which depends on your audit trail, your monitoring, and your appetite for being wrong in public. It does not depend on what the model can do.

A meeting-scheduling agent reaches level four in a week. A multimillion-dirham procurement decision should not.

> **Autonomy is earned, not granted.**

Five questions worth answering this year

Most large organisations have an AI operating model by now. A team, a governance forum, a data platform, a pipeline of use cases. None of it wasted. But it was built for AI that recommends, and it does not stretch to AI that acts.

What comes next needs answers to five questions. Call it an Autonomy Operating Model.

01 **What work should machines perform?** Redesign whole workflows, not individual tasks. McKinsey Global Institute puts close to $2.9 trillion of US value on the table by 2030, but only for organisations that rebuild work around people, agents and robots together, rather than automating tasks inside a process designed for humans alone.

02 **How much autonomy should each machine hold?** The ladder above, applied case by case, and revisited as the evidence comes in.

03 **Who governs machine authority?** Every agent needs the equivalent of a passport: identity, role, permissions, authority, limits, escalation, audit trail, owner.

04 **How do we plan across three kinds of capacity?** Human, digital and physical. “How many people do I need” becomes “what mix gets this done, and who is accountable for each part of it.”

05 **What are we actually measuring?** Cost per completed outcome, and how often a human has to step in. Not copilots deployed.

Everyone will have the same models

Models will keep improving. Robots will get cheaper. Access to intelligence is commoditising quickly, so a model is not going to be anyone’s advantage for long.

What is hard to copy is the organisation around it. Knowing what to automate and what to leave alone. Knowing which decisions stay with people because they should, not because the technology has not caught up yet. Knowing how much authority a machine can hold in your particular risk environment, and being able to demonstrate that afterwards to a regulator, a customer, or a board.

The question stops being whether an agent is intelligent. It becomes whether it is authorised.

OVER TO YOU

What is the highest rung on the Autonomy Ladder your organisation could genuinely govern today, and what is actually holding it there?

Sources

Every figure in this piece was checked against the primary publication. Dates are given so readers can see how current each claim is.

1. **BMW Group**, BMW Group to deploy humanoid robots in production in Germany for the first time (2026). Spartanburg pilot results, the Leipzig pilot, and the Center of Competence for Physical AI in Production.
   press.bmwgroup.com/global/article/detail/T0455864EN
2. **Figure AI**, Figure 02 BMW deployment results and fleet retirement (November 2025). 90,000+ parts, 1,250+ operating hours, 30,000+ vehicles.
3. **NIST / NCCoE**, Accelerating the Adoption of Software and AI Agent Identity and Authorization, draft concept paper (5 February 2026).
   csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd
4. **McKinsey Global Institute**, Agents, Robots, and Us: Skill Partnerships in the Age of AI (November 2025). ~$2.9tn of US value by 2030, conditional on workflow redesign.
5. **Microsoft**, 2026 Work Trend Index Annual Report (5 May 2026). 67% versus 32% organisational-to-individual impact.
   microsoft.com/en-us/worklab/work-trend-index