



# Agentic AI, done wrong, can be a house built on sand

David Norfolk ©Bloor Research July 2026

**This blog was prompted by an interesting article in [Futurism](#). According to some estimates, it reports, up to 79 percent of US corporate execs have some type of AI agent in the making – but one Gartner prediction found 40 percent of these projects will implode due to poor risk controls (see Gartner citation below).**

Futurism reports an “anti-pattern” (a common way of doing something that predictably delivers a bad result) identified by Sayali Patil in [VentureBeat](#). She documents a case in which an Agentic AI shut down a server in response to an emerging issue, which then caused unconnected but mission critical services also running on that server to fail: “The blast radius of that agent action was not the service restart. It was everything downstream of the restart, in a system state the agent had no complete picture of”, Patil writes.

Nevertheless, this really isn't a new AI issue. It is the sort of issue that has always arisen, even before AI became a thing, when concentration on fixing an immediate problem resulted in unanticipated, and undesirable, consequences elsewhere. I remember once, from my days in DBA (Database Administration), when the deadline for a new banking system was so important that all its compile jobs took priority over everything else – even over the jobs DBA was running (or trying to) in order to implement the necessary database changes for the new system. The result? Total deadlock and not an AI in sight. If you abstract away the technology, this is the same anti-pattern as in our AI example: involving poor system design, poor analysis of the scope of impact, tunnel vision focused on just one issue and, presumably, no testing (simulation, role-playing) of the possible impact of a decision, whether made by a stressed human or by an Agentic AI responding simplistically to an obvious issue. As an aside, might the Agentic AI learn from this failure and do better next time? I suppose that that could be arranged, but we've already had an expensive business outage and one such is one too many.

I asked Bloor's Chief Analyst Richard Skellet for his view on this, and he confirmed that “it's not really AI agents botching tasks, it's a technology immaturity problem, not an implementation problem. It is the direct and predictable consequence of deploying Agentic capability on top of a data and architecture foundation that was never designed to support it”.

Recognition of the inadequacy of the foundations for Agentic AI shouldn't come from hindsight; the MAIM (Mainframe Application Intelligence/Modernization) diagnostic, part of Bloor's [Fusionwork](#) Methodology, for example, should make the issue visible before deployment. Richard says: “A rogue agent had execution rights without environmental awareness. It could action a task. It could not model the system state that task would disturb. That gap between what the agent is permitted to do and what it is capable of understanding is not closed by better guardrails or improved change management after the fact. It is closed, or not closed, at the architectural design stage”. The [MAIM](#) program is Bloor's current work-effort around the evolution of Master Data Management with the addition of a new data intelligence model.

What is often missed, not just in articles like the one prompting this blog but also in much analyst advice, is the reason for the Agentic agent not having a complete picture of the system state. This is, fundamentally, because the data layer underneath it was never unified. MDM, integration, governance, quality – the multi-billion dollar vendor landscape of fragmented point solutions – exists precisely because integrating architectures are routinely neglected. This means that the agent in this anti-pattern was not just architecturally under-powered, it was sitting on top of a data foundation that could not have given it a complete picture of system state even if the agent had been designed to ask for one.

Bloor's position is that this is the [MAIM](#) argument made concrete. The cascade failure is not an Agentic AI failure. It is a “Hierarchy Tax” ►

ARTICLE REPRINT



- ▶ (that is, it is the productivity loss, innovation bottlenecks, and hidden costs associated with rigid, top-down management structures), which has always been there potentially, but which is now made visible by new Agentic loads. In other words, it is an emerging structural liability, compounded from technical debt, digital transformation stall, and AI debt, which is now made manifest as a live operational incident with an impact “blast radius” encompassing much of the organization.

Gartner frames its 40 percent Agentic AI project collapse prediction as a **risk controls** issue but poor risk management is the symptom rather than the cause. The risk controls have nothing structurally sound to govern. GORC – Governance, Oversight, Risk and Compliance – cannot function as a post-deployment audit layer on top of an architecture that was never designed for governed Agentic operation. GORC is a pre-condition of safe deployment, not a response to its failure. The organization suffering the “Agentic failure” we are talking about did not have a GORC architecture. It only had a change management plan and a vendor contract.

The **MOM** (Mutable Operating Model) layer makes the same point from the operating model point-of-view. Agentic AI is not simply a new interface. It is a redefinition of the role of humans and machines in the workflow. Richard Skellett tells me that he has seen Agentic AI deliver productivity gains, but only where the operating model was designed to absorb Agentic capability. Other organizations, such that described in the Futurism article that prompted this blog, organizations deploying task execution into an unmodelled, and therefore poorly understood, environment and calling it transformation, are not achieving the potential benefits

The CEO dimension compounds the problem at the decision layer. Agentic AI is at the top of the hype cycle right now and many managers see it as a new “silver bullet” that will kill all of their problems. Executive expectations are calibrated against an AGI-era framing (against an era of Automated General Intelligence – with truly “thinking machines”) – and pre-suppose the use of contextually aware, self-correcting technology, reasoning across system states. Production Agentic AI is currently hitting the gap between bounded task execution and incomplete whole-system – environmental – models. That gap does not close with better vendor briefings. It closes with architectural honesty before the program is approved; and that requires a diagnostic instrument — which is partly what Bloor Bullseye (a model that examines the gap between what technology does and what a perfect product would do) is designed to be.

The fundamental Bloor position, according to Richard, is this:

*“...you cannot safely deploy Agentic AI on top of fragmented, unmanaged data. You cannot govern what you have not architecturally designed for governance. And you cannot expect an agent to model and manage blast radius if the data layer beneath it has never been unified enough to define what the system state actually is”.*

ARTICLE REPRINT



**BLOOR  
RESEARCH**

+44 (0)1494 311 460  
info@Bloorresearch.com  
www.Bloorresearch.com