ESSAY

The Artefactual Organisation

As agentic AI becomes the adaptive architecture of the organisation, the critical boundary is not how much it does, but who decides what matters.

Etienne van der Walt, MD

Neurologist | Living Systems Thinker | Founder, Neurozone | Author of 'To Emerge'

16 October 2026

From Companion to Organisation

The conversation about artificial intelligence at work is moving rapidly beyond automation. Generative AI already participates in human thinking: we use it to explore, challenge, deliberate, analyse and create. In this relationship, AI functions as a companion. The human still determines what matters, brings context and judgement, and remains responsible for the decision. Companion AI thinks with us.

Agentic AI changes the relationship. Agents do not merely participate in cognition; they act. They can be given an outcome, undertake a sequence of tasks, use tools, coordinate with other agents and return with completed work. Increasingly, organisations will contain layers of agents performing research, coding, analysis, monitoring, administration and commercial activity, with some agents coordinating the work of others.

Yet even a highly agentic organisation can remain fundamentally human. Human executives determine direction. Human managers own functions. Human teams remain responsible for product, commercial activity, finance, legal and operations. Artificial agents make those processes faster and more capable, but the organisation is still constituted around human agency.

There may, however, be another organisational form beyond this. I want to propose a term for it: the *artefactual organisation*.

The distinction is not primarily about how many humans and how many artificial agents an organisation contains. It is about where the adaptive agency of the organisation resides. The important transition occurs when artificial systems no longer merely support the organisation's principal processes, but increasingly constitute its capacity to sense, decide, act, learn and adapt.

The Functional Flip

Product development provides a useful example. Initially, AI may help a product team summarise customer feedback, analyse usage data or write code. As agentic capability develops, interconnected agents could continuously analyse product implementation, customer behaviour and wider market signals. They identify emerging needs, predict opportunities and refine the product backlog. Product-development agents translate those requirements into specifications, build functionality, test it, detect failures, correct them and prepare the next release.

At first a human remains inside the loop. A product owner approves the change. But the human can progressively move outward. Instead of managing the process, the product owner receives a report describing what changed in customer behaviour, which opportunity the system identified, how the backlog was reprioritised, what was built, what testing demonstrated and what the system intends to pursue next.

At some point, artificial intelligence is no longer simply assisting product development. Product development itself has become artefactual.

Now extend the same architecture into the market. Customers use the product and generate outcomes, behaviour and feedback. Artefactual systems continuously sense those signals and combine them with broader commercial information. They identify new customer needs and potential markets, develop propositions, test pricing, evaluate channels and refine commercial approaches. Those findings feed back into product development, while new product capabilities create new commercial possibilities.

A continuous adaptive loop begins to form. Customer behaviour changes the organisation's understanding of the market. Market understanding changes the commercial opportunity. The opportunity changes the product requirement. The product changes. Customers respond to the changed product, and their response becomes the next signal. The organisation continuously senses, predicts, builds, tests, sells, learns and adapts.

Finance, legal, compliance and administration do not disappear. They increasingly become supporting artefactual functions around this central product-to-market loop.

This produces an important inversion. Today we ask where AI can augment humans. An artefactual organisation could increasingly ask where a human would improve the performance of the system. If a particular market requires embodied trust, negotiation or human relationship-building, the system may use people. If another market can be served more effectively without them, it may not.

Human labour therefore becomes one capability available to the organisation rather than the substance from which the organisation must necessarily be constructed.

This is the functional flip: the point at which humans are no longer operationally necessary inside the principal adaptive loops through which the organisation senses its environment, develops its products, engages its market and responds to what it learns. In this arrangement, humans remain valuable, even essential in particular contexts, but the organisation no longer requires their continuous participation in order to operate and adapt.

The functional flip does not remove humans from the organisation altogether. It moves them. As in every earlier wave of automation, the people who leave the operating loop do not vanish from the system; they move above it - into governance, and into the maintenance work of calibrating, auditing and correcting the systems that now run below them. The real question is whether that layer above remains capable of genuine oversight, or slowly becomes a layer of signatures.

> *The functional flip is not defined by how many agents an organisation contains. It is defined by where its adaptive agency resides - and by whether the humans who move above the loop can still genuinely govern what runs beneath it.*

When Strategy Begins to Emerge

We tend to assume that even if AI performs much of the operational work, strategy will remain human. But the adaptive loop complicates that assumption.

Customer behaviour produces evidence. Evidence changes predictions. Predictions reveal opportunities. Commercial experiments test those opportunities. Their results alter product priorities. Product development creates new capabilities. Those capabilities create new markets, and market responses generate new evidence.

What we conventionally call strategy can begin to emerge from the interaction among these processes.

This reverses the familiar organisational hierarchy in which vision produces strategy, strategy determines operations and operations engage the market. In the artefactual organisation, information flows powerfully in the opposite direction. Product refinement follows business need. Business need follows what implementation and customer behaviour reveal. Commercial strategy responds to what the market is showing, and commercial

predictions feed back into product development. Product and market continuously refine one another.

Strategy becomes less a periodic plan created by senior executives and pushed downward through an organisation, and more an emergent property of the organisation's continuous interaction with its environment.

Humans may initially evaluate the strategic options generated by the system. Over time, they may increasingly sign off on strategies that the artefactual organisation itself has developed. Eventually, routine strategic approval may no longer add sufficient value to justify remaining inside the loop.

The human role has moved again. Humans provide vision. The artefactual organisation increasingly provides strategy and execution.

Emergence Relative to What?

There is, however, a boundary to this argument. An adaptive system cannot simply adapt in the abstract. Strategy can only emerge relative to something. The system has to distinguish signal from noise, determine what constitutes success and decide which outcomes should be amplified or avoided.

Revenue may matter. Customer retention may matter. Margin, market share, human wellbeing or environmental sustainability may matter. But something has to establish what matters.

This is where salience becomes decisive.

In living systems, salience is the directional capacity to recognise what matters and to allocate attention, energy and resources accordingly. It is distinct from capability. An organism, a team or an organisation can possess enormous capability while directing it toward the wrong objective.

The same holds at organisational scale. An organisation can become highly coherent around something that no longer matters. It can operate with extraordinary speed and efficiency while becoming progressively better at producing the wrong outcome.

Vision therefore performs a salience function. It establishes the field within which adaptation has meaning. It tells the organisation which signals deserve attention, what

constitutes progress, what should be optimised and, importantly, what should not be sacrificed in pursuit of an objective.

This suggests a useful distinction. The functional flip occurs when the system no longer needs humans to determine how. A deeper transition occurs if it no longer needs humans to determine what matters.

> *An adaptive system can become extraordinarily good at pursuing an objective. But successful adaptation does not, by itself, establish whether that objective is worth pursuing. Vision is not a statement on a wall. It is the field of salience within which adaptation acquires meaning.*

The Human Flip

Once operations and strategy become increasingly artefactual, it becomes difficult to avoid the next question. Why should vision necessarily remain biological?

An artefactual system capable of sensing its environment, comparing possible futures, allocating resources, developing strategies and evaluating outcomes can increasingly reproduce functions currently associated with executive leadership. It can ask what is changing, which opportunities are emerging, what should be built, what should be stopped, where capital should be allocated and which capabilities are required.

But one question is categorically different: what should this organisation be trying to achieve?

The human flip would occur when humans no longer determine the field of salience within which the organisation's adaptive loops operate.

There is a reason to treat this question differently, and it is biological rather than sentimental. Artificial intelligence operates powerfully at the symbolic and representational layers of intelligence - language, pattern and structured reasoning. There, it can reason about objectives, compare them and weigh options against them with remarkable precision. But in living systems, the sense of what matters does not originate in those layers alone. It rises through the full depth of an organism that has something at stake: a body that can be sustained or harmed, relationships that can be strengthened or broken, a future it will have to live in. An artefactual system can compute importance relative to the objectives it has been given. It does not bear the consequences of those objectives. The people inside and around the organisation do.

Salience is therefore not simply one more function to be automated once systems become capable enough. It is the point at which an organisation remains anchored to the living systems it exists to serve.

Importantly, the human flip need not happen through a dramatic transfer of control. It could happen through drift. An organisation begins with a human purpose, such as serving its customers well. Retention becomes a useful proxy for whether customers are being served, so the system optimises retention. It then identifies the signals that best predict retention and optimises those. Each step is rational. Each step appears to improve performance. Yet gradually the proxy can become the purpose. This is a familiar pattern in any measured system - when a measure becomes a target, it tends to stop being a good measure - but an artefactual organisation could run that cycle continuously, at a speed and scale no human institution has ever done.

The formal vision may still belong to humans while the effective vision has migrated into what the system repeatedly measures, rewards and reinforces.

Catching that drift is precisely the work of the human layer that remains above the artefactual loop. But it can only be done by people who still hold salience themselves - who know, in an embodied sense, what the organisation is for, rather than only what its dashboards report. Governance exercised by humans who have drifted into approving whatever the system recommends is not genuine governance. It is the human flip, already under way.

The Same Pattern at a Smaller Scale

There is a more familiar place where this pattern already plays out. Anyone who works closely with an AI companion knows the division of labour, explored in detail in the framework essay Hybrid Intelligence in Organisations. The companion brings analytical reach and generative capacity. The human brings the embodied sense of what matters, and decides what the collaboration is for.

That relationship also evolves. Early on, the human directs closely - step by step, section by section, checking each output. As trust is earned, the human hands over more of the generative work and moves to reviewing whether the result serves the intent. This is not laziness. It is the functional flip happening inside the dyad: the human leaves the loop of how, retains the question of what, and remains accountable for the outcome.

The dyad and the organisation are not the same, and the differences matter. In the dyad, distance is easily reversed; the human can step back in at any moment. In the organisation, salience is exercised further away, through objectives and measures rather than direct review. And an organisation's salience cannot belong to a single person. It has to be held collectively, which depends on the coherence of the humans who hold it.

Yet the pattern holds, and so does the boundary. The risk in the dyad is not distance. It is that the companion's framings quietly become the human's - not through outputs that are wrong, but through outputs that are good enough - until the human can no longer tell whether the work still serves what matters. That is the human flip in miniature. The safeguard is the same at both scales: humans who remain accountable, who periodically step back in, and whose embodied judgement stays sharp enough to recognise drift.

> *The human flip is not the moment the human stops checking every step. It is the moment the human can no longer tell whether the output serves what matters.*

Two Flips, Two Kinds of Stakes

This is why the distinction between the two flips matters. The functional flip is principally a question of organisational design and performance: where do humans improve the outcome, and where do they no longer need to remain in the loop? It can be measured, governed and, in principle, reversed. The human flip is a question of purpose - and it can be crossed without anyone ever deciding to cross it.

Every individual delegation may be rational. At every stage we can tell ourselves that humans remain in control because we have simply moved upward to the next level of supervision. But supervising operations is not the same as determining strategy. Approving strategy is not the same as creating it. And approving what a system recommends is not necessarily the same as determining what that system should ultimately be for.

The artefactual organisation is therefore worth considering before it fully exists. The important question raised by increasingly agentic AI may not ultimately be how many jobs it replaces, how much productivity it creates or even how autonomous organisations become. It may be about the boundary we choose to maintain between agency and salience.

> *We may become comfortable allowing artefactual systems to become extraordinarily capable at determining how. We should be far more deliberate before allowing them to determine what ultimately matters.*

Because underneath every adaptive system, every strategy and every optimisation problem sits a question that cannot simply be optimised away:

What is the system for?

About the Author

Dr Etienne van der Walt, MD, is a neurologist, entrepreneur, global keynote speaker and living systems thinker whose work sits at the intersection of neuroscience, biology and leadership. As the founder of Neurozone, he has spent more than two decades developing evidence-informed frameworks that help leaders and organisations improve resilience, wellbeing and sustainable performance. His work challenges traditional mechanistic approaches to leadership by viewing organisations as living systems that require intelligent biological regulation rather than increasing control. He is the author of To Emerge and writes and speaks internationally on resilience, leadership, organisational health and the science of human flourishing.

Etienne is currently on a quest to operationalise the yield framework in the age of hybrid intelligence - to develop the measurement tools, diagnostic frameworks, and practical methodologies that allow leaders and organisations to design, monitor, and sustain biological yield as human and agentic AI systems deepen and evolve. He welcomes conversation and collaboration with those exploring the same frontier.

The living systems account of salience and adaptation on which this essay draws is developed in full in To Emerge: Leading a Living System, From Cells to Society.

This essay builds on Hybrid Intelligence in Organisations, the framework essay that sets out the companion, agent and maintenance layers referred to here, and the biological architecture of salience on which this argument rests.

For more essays and resources, visit www.etiennevanderwalt.com

Selected Reading

Bainbridge, L. (1983) 'Ironies of automation', Automatica, 19(6), pp. 775–779.

Capra, F. and Luisi, P.L. (2014) The Systems View of Life: A Unifying Vision. Cambridge: Cambridge University Press.

Damasio, A. (2010) Self Comes to Mind: Constructing the Conscious Brain. New York: Pantheon Books.

Harari, Y.N. (2024) Nexus: A Brief History of Information Networks from the Stone Age to AI. New York: Random House.

Kauffman, S. (1995) At Home in the Universe: The Search for the Laws of Self-Organization and Complexity. New York: Oxford University Press.

Parasuraman, R. and Manzey, D.H. (2010) 'Complacency and bias in human use of automation', Human Factors, 52(3), pp. 381–410.

Russell, S. (2019) Human Compatible: Artificial Intelligence and the Problem of Control. New York: Viking.

van der Walt, E. (2026) 'Hybrid Intelligence in Organisations', Framework Essay. Available at: https://etiennevanderwalt.com/wp-content/uploads/2026/07/Hybrid_Intelligence_Framework_Essay-Etienne-van-der-Walt.pdf