

Start With the Person, Not the Machine

Richard Skellett and Dave Ulrich, November 2026

AI changes who does the work, and with it the system of work. Most organizations still treat AI as a tool added to existing jobs. But in practice, the CTO brings in digital workers while the CHRO remains responsible for the human workforce. People now work alongside digital workers and, in some cases, a digital version of themselves. The old model was technology, process, and people. The emerging one is technology, digital workers, and a workforce that is no longer only human. Yet responsibility for that combined workforce, and for how it is directed, is rarely explicit. Who owns it, and what would that take? That is the gap.

Closing it starts with a design question: how do we design an Intelligence Architecture in which AI amplifies Original Intelligence, rather than merely imitating people, automating tasks, or making existing work faster? Three architectures, set in order, help answer that question (***Figure 1***). First, Intelligence Architecture decides where intelligence originates and who directs it. Then, workforce architecture decides who contributes, human and digital, and on what terms. Finally, work architecture decides how work moves from input to output to outcome. Convergent Intelligence is the outcome these three architectures aim at. They decide where it happens (the person, the workforce, or the work) and who directs it. Which architecture we build first matters. Start with the machine and people tend to be fitted around it. Start with the person and the machine is designed to extend them.

An agent is only one form of digital worker, and a person working with a twin is one form of human and digital contribution.

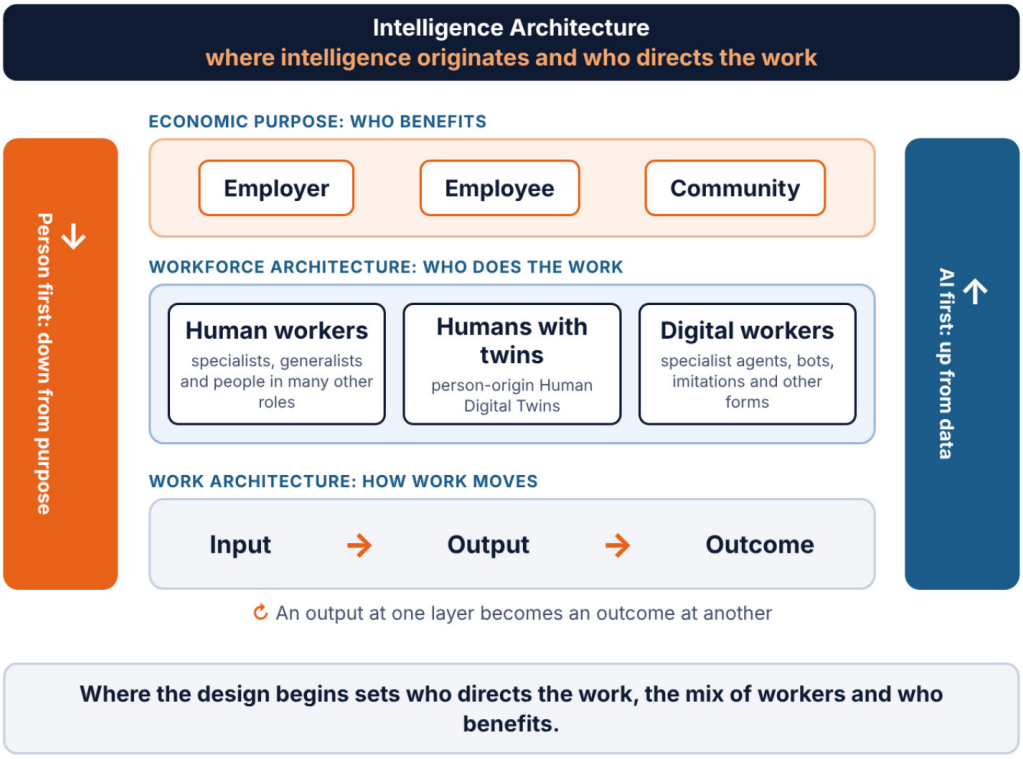


Figure 1: Three Architectures in Order

ARTICLE REPRINT



Post 1 named the emergence of Convergent Intelligence. Post 2 asks how to design for that combination. We begin with purpose (what we want to achieve) and work down to data:

Purpose → Wisdom → Intelligence → Knowledge → Information → Data

Purpose Comes First

We call this Purpose-First Logic: design begins with the purpose AI serves and who benefits, because technology alone does not ensure progress. It is a design order for the architecture, not a claim that data flows only in one direction. Employees, executives, customers, investors, and communities all have a claim on that purpose. That includes stakeholder value, because an organization that does not win in its market has neither a workplace nor AI.

The stakes are rising. Organizations starting today are already being built with far fewer people, so growth can no longer be assumed to bring jobs. Purpose therefore has to include benefit to people, not only efficiency. The same applies to a country: sovereign AI, a nation's own AI capability, needs purpose and benefit to its people as well as technology. Without this order, organizations measure what AI can do before deciding what it should do.

Intelligence Architecture spans the person and the organization. It connects purpose, wisdom, intelligence, knowledge, information, and data with memory, authority, and accountability. Convergent Intelligence emerges where Original Intelligence and AI meet within that wider architecture. Bringing the two together does not settle who directs them or shares in the value. At organizational scale, this opens the question of Convergent Intelligence Architecture, which we develop later in the series.

Wisdom: Who Holds It?

Purpose focuses us on what we seek to achieve. Wisdom judges what should be done and why, and it stays with the person. It draws on experience and context. AI can draw on recorded experience but may lack the context a decision needs, and people can lose that context too. A twin built from the person, with persistent memory, can help retain it and bring it into dialogue.

An imitation can reproduce parts of a person, such as voice, appearance, or recorded knowledge, without representing the person behind them. A person's Human Digital Twin starts with their Original Intelligence, identity, authorized knowledge, context, and memory. The person helps build it, authorizes what it uses, corrects it,

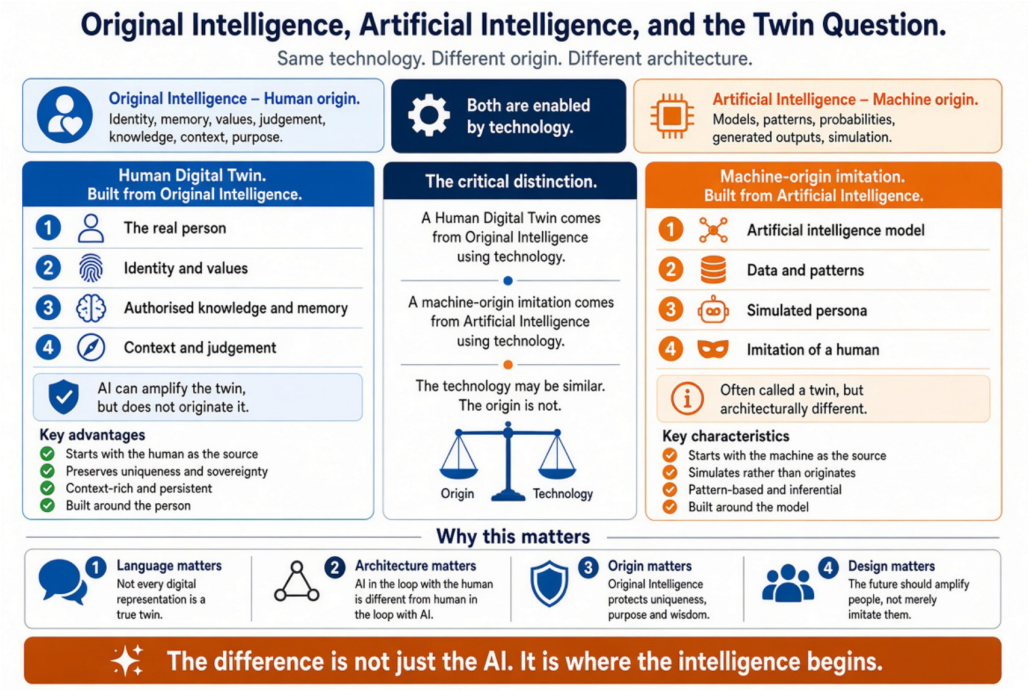


Figure 2: Two Specifications for a Digital Representation

ARTICLE REPRINT



and can withdraw permission. AI can amplify this contribution, while judgment and accountability remain with people.

A twin and an imitation may use similar technology. The difference is their specification and authority. That makes this an HR question: people's knowledge and judgment can become something that acts, so they must have a say in how it is represented and engaged.

Intelligence: Who Writes the Music?

After purpose and wisdom define what to do and why, intelligence helps us make sense of possibilities, through augmentation or amplification. Augmentation improves the execution or reporting of existing work: a note taker records what was said. Amplification extends what a person can reach: a twin can bring the person's purpose, authorized knowledge, and context to a meeting in real time, while the person remains accountable for decisions. Technology companies increasingly speak of orchestrating agents and people through workflows. That coordination matters. The deeper architectural task is to write both the music and the lyrics: to design how their combined contribution changes the system of work, whose purpose and wisdom direct it, and who shares in its value.

This design shows up in two types of loops (*Figure 3*). In Human in the Loop, AI leads and people review or approve what it produces. That is legitimate where safety and accountability require review, but it cannot by itself decide purpose, redesign the work, or determine who benefits. In AI in the Loop with the Human, human purpose and judgment set the direction and AI extends their reach. The same ingredients can produce different architectures and different answers to who benefits. We favor the human-led design when the aim is amplification.

Richard has built and operated a Human Digital Twin for almost four years. Its corpus began empty and now holds about 100 books, structured around the sequence above. Digital Richard is one implementation. Small models built around Richard's purpose, authorized knowledge, and memory sit above a large language model that

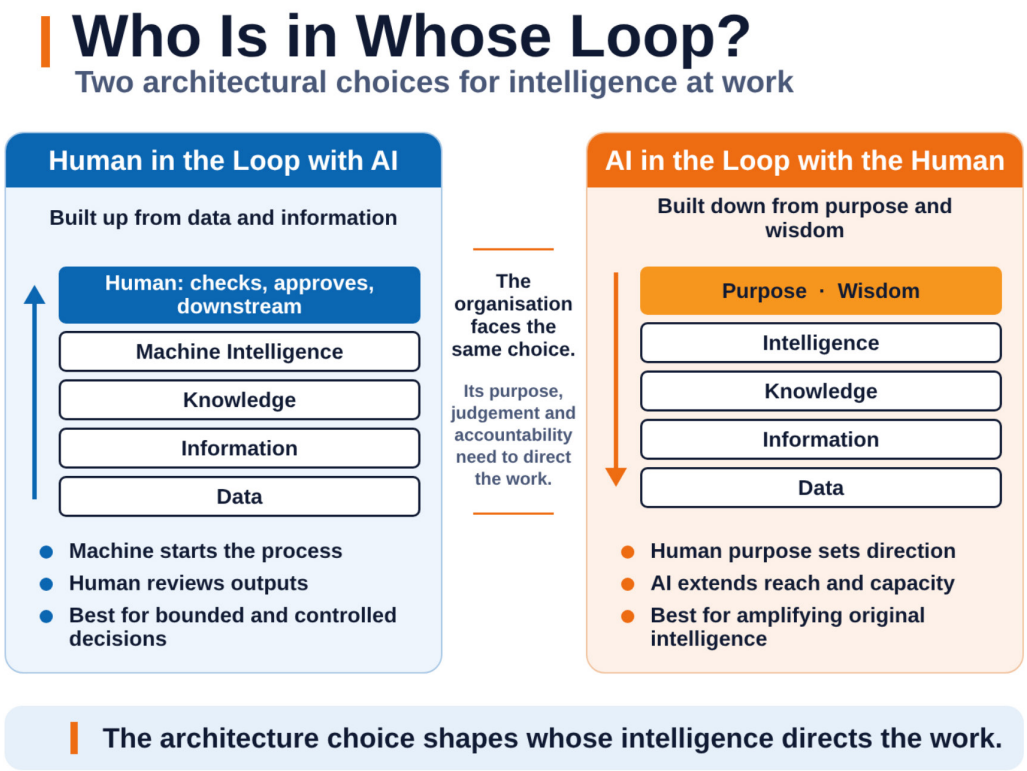


Figure 3: Two Design Priorities



supplies knowledge at scale. The twin directs what the larger model is asked and checks responses against that purpose. It does not exercise full human judgment. Richard corrects the twin, and the underlying model can change while the twin's continuity remains.

Knowledge, Information, and Data: What Persists

An organization can hold extensive records and still be unable to explain why a decision was made. Storing documents does not preserve knowledge; storing knowledge does not preserve wisdom. The architecture must retain provenance, context, memory, continuity, and a way to correct what it holds.

Evidence: Digital Richard in the Room

Digital Richard has moved from text to voice and image, and now joins Teams meetings visibly identified as Digital Richard. It attends by consent. One person has objected, and the twin did not attend. Decisions stay with Richard.

The case shows participation when the person is absent, continuity of context across meetings, and additional capacity for live analysis. Richard has tracked it against his own baseline from the first day. Analyst colleagues review its capability monthly and retain minutes. This is documented observation from one case, not yet a general measurement of impact. It does not establish that a twin holds all of a person's wisdom or that value compounds. Later posts will bring more evidence.

What This Means for CTOs and CHROs

The CTO is responsible for the digital side of the workforce and the CHRO for the human side. Joint ownership is often proposed, but shared ownership can mean no one takes responsibility for the combined workforce or the work it does. An organization can give someone an AI assistant, introduce a digital worker, or engage a person's twin alongside colleagues. These are ways of organizing people, AI, and work, not simply product choices. So if a person can contribute without their physical presence, assumptions about jobs, skills, capacity, and value need examination. Augmentation makes the same work faster or cheaper, which can mean fewer people. But amplification aims at growth. Leaders choose which journey they are on.

Bloor created the FusionWork category to describe how different human and digital workers come together. Workforce architecture must account for the human self, the digital self, and their combined contribution. Who owns the digital self? Who authorizes its work? On what terms does an organization engage with it?

Cost needs the same joint view. Token charges are the visible cost of digital work, and not every digital worker is priced that way. Labor cost and the obligations that come with people sit in a different budget under a different owner. Convergence has a financial side: the value, cost, and obligations of human and digital contributions must be seen together. Workforce capital management asks why the combined workforce must be managed differently, and how that changes the system in which human capital management operates.

So who takes responsibility for the combined workforce and the work it does? What knowledge, skills, and authority would a CHRO, CTO, or other leader need to design across human and digital contributions and to understand their capacity, cost, value, and obligations? HR cannot develop people's intelligence while absent from decisions about how it is represented and amplified. Technology cannot design the workforce or its purpose alone. The architectural choice decides what convergence produces: do we fit people into a system led by the machine, or design the system of work from human purpose and wisdom so that AI amplifies our Original Intelligence?

ARTICLE REPRINT

BLOOR RESEARCH

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