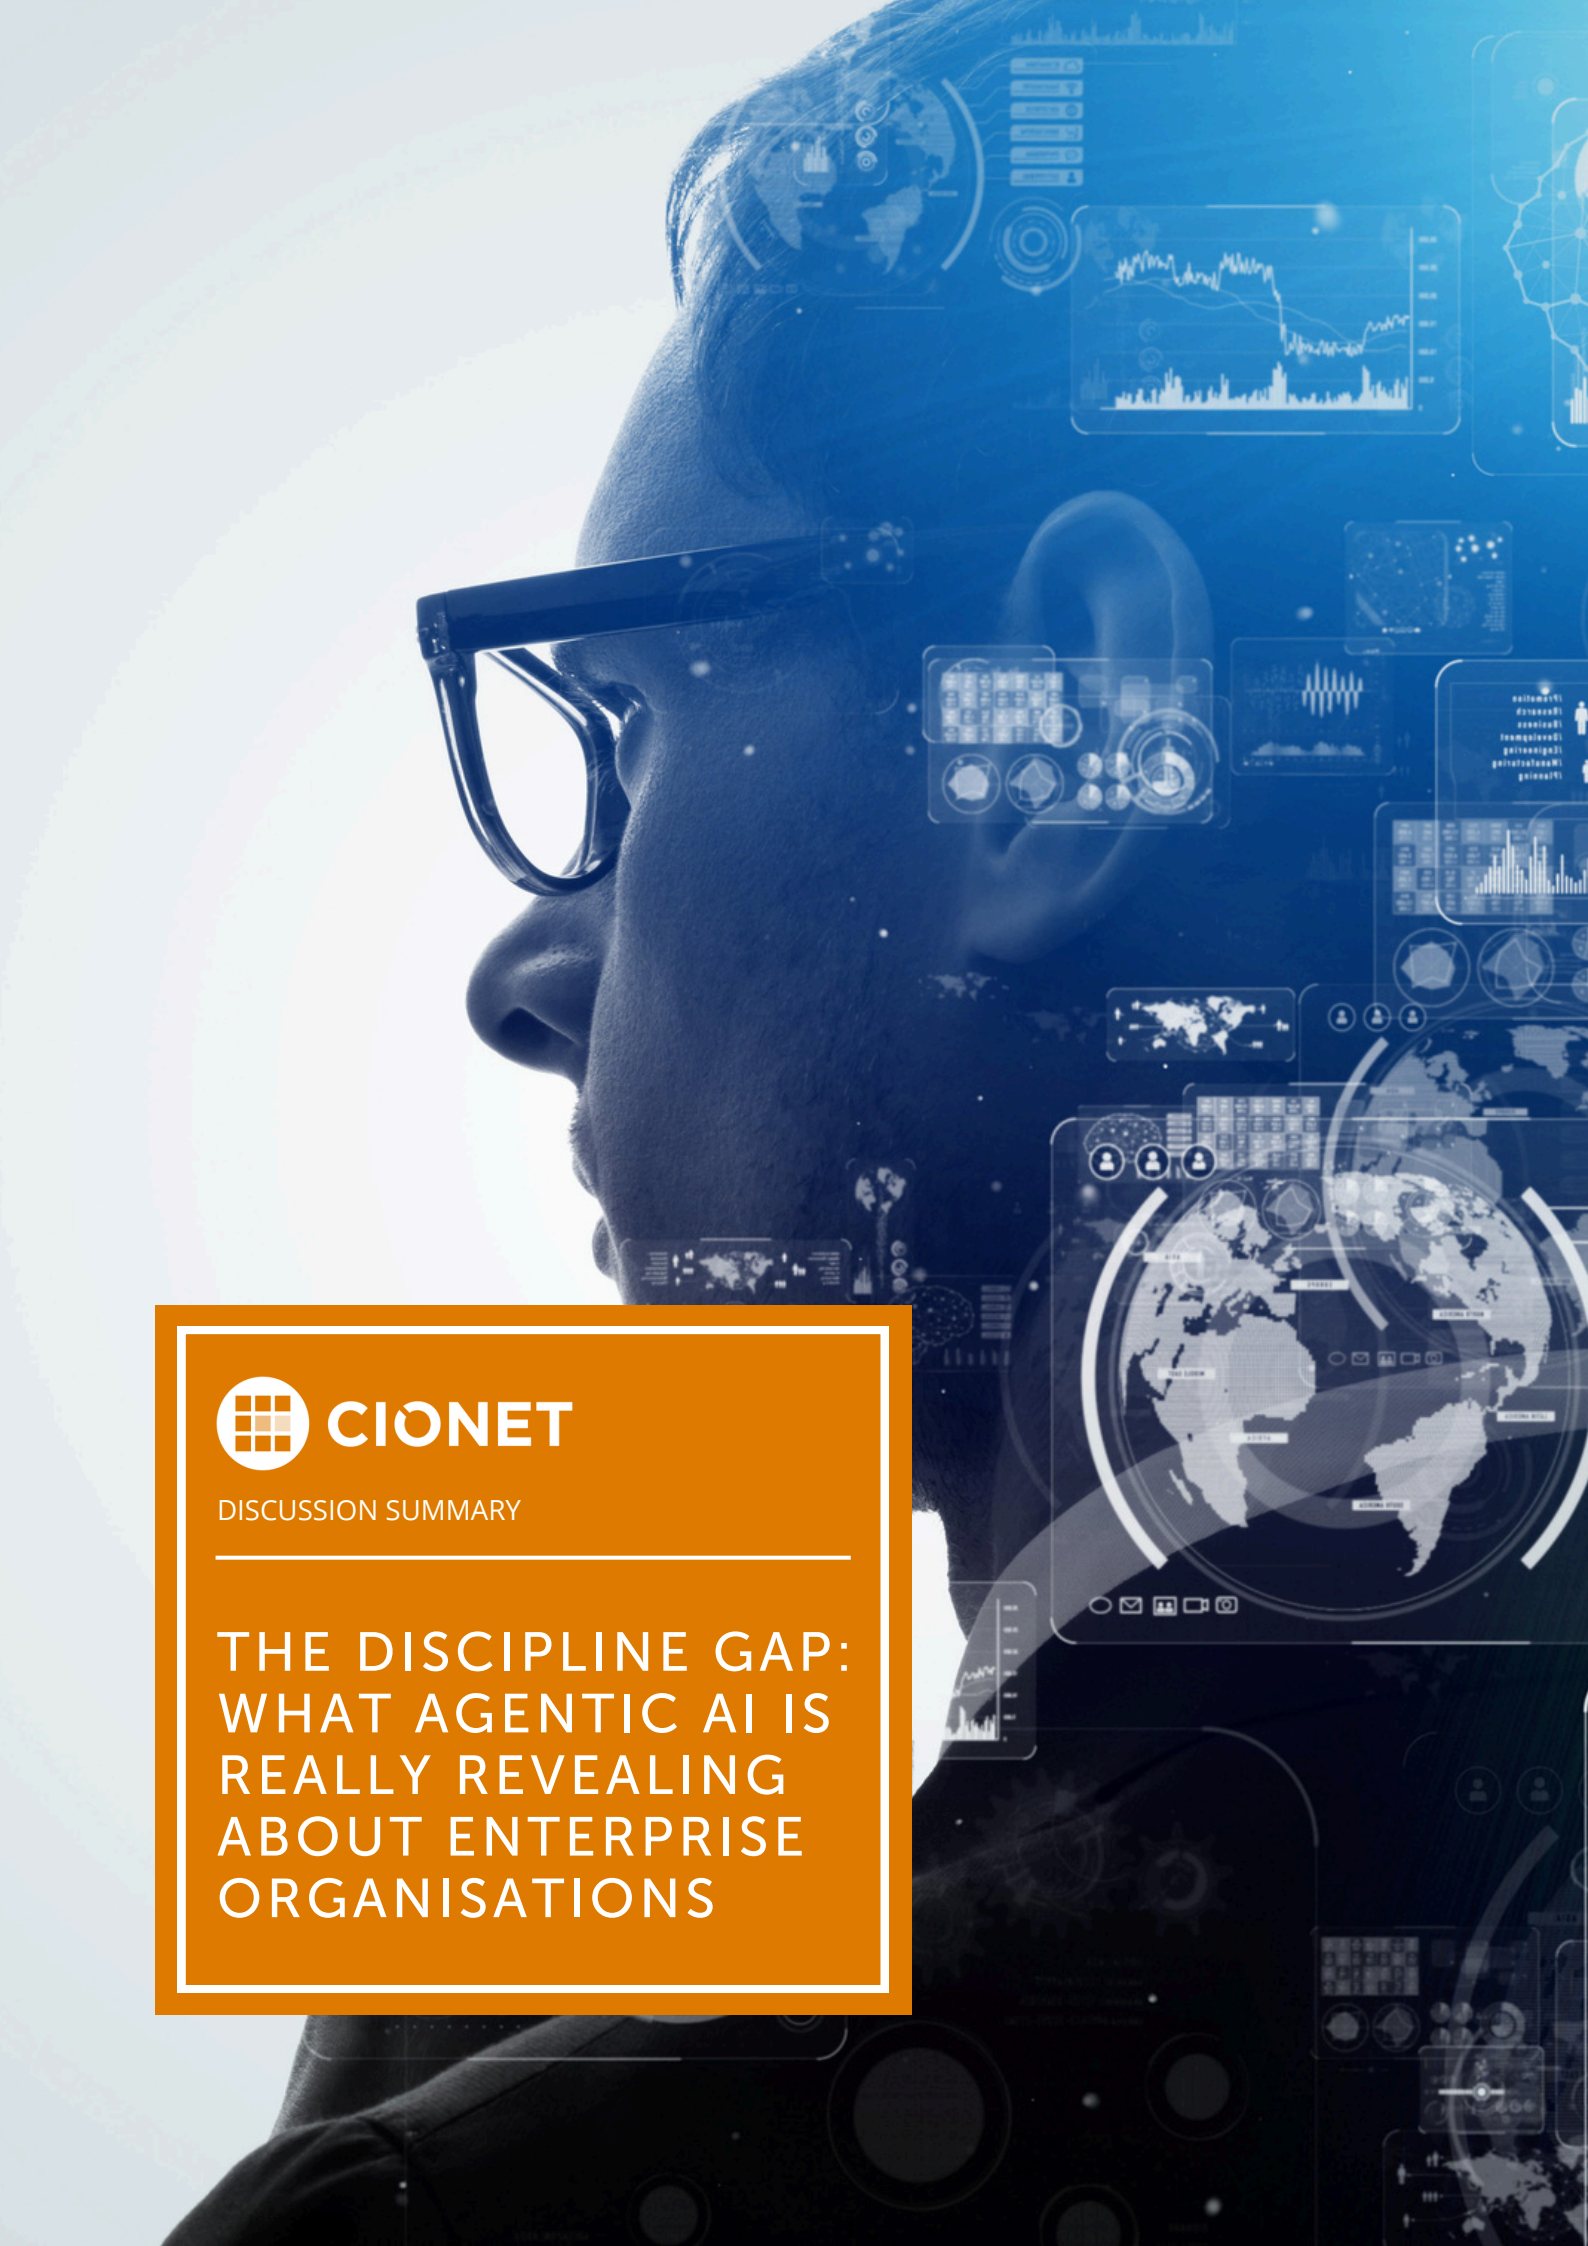

CIONET

DISCUSSION SUMMARY

THE DISCIPLINE GAP: WHAT AGENTIC AI IS REALLY REVEALING ABOUT ENTERPRISE ORGANISATIONS

The Discipline Gap: What Agentic AI Is Really Revealing About Enterprise Organisations

The conversation technology leaders are having about agentic AI is not quite the one they think they are having. The surface debates concern governance frameworks, talent acquisition, and adoption pace. What keeps surfacing underneath is more awkward: a gap between what organisations can picture and what they can actually deliver. That gap has less to do with technology than most of the discussion assumes.

The Governance Illusion

There is a version of the agentic AI story where governance is the problem, the institutional drag that slows down people who would otherwise move faster. The leaders working through this in practice tend to describe something different. Governance done badly is the risk. Governance done well is what makes deployment possible at all.

The practical reality is illustrated clearly by one participant's experience. A customer-facing agent was technically complete and ready within weeks. What followed was not deployment. It was over a year of navigating regional regulatory variation, legal sign-off, and internal consensus across marketing, sales, and risk. The technology was the straightforward part.

This will not surprise anyone who has worked inside a large enterprise for any length of time. Shared infrastructure, heterogeneous systems, and organisational process mean adoption curves rarely match vendor timelines. The previous wave of digital transformation is, in many organisations, still ongoing. There is no particular reason to expect agentic AI to behave differently.

Accountability for agent errors remains genuinely unresolved. One framing that found some agreement: the creator of the agent carries the responsibility, in the same way an organisation is accountable for the processes it puts in place.

That holds reasonably well for contained use cases. It becomes harder to apply in multi-layered systems, where a decision made fifteen steps upstream in an agent chain may be almost impossible to trace back. One participant made the point that provability (being able to show what an agent did and why) is a more pressing challenge than the accountability question itself. Establishing who is responsible matters less if the chain of decisions cannot be reconstructed.

The Talent Arithmetic That Does Not Add Up

The workforce debate around AI tends toward extremes. Either jobs disappear rapidly or the concern is overblown. The picture that emerged from leaders operating at scale is more specific, and more demanding, than either position suggests.

There is a genuine shortage of people who have built agentic systems at enterprise scale rather than in controlled environments. What functions well in a pilot regularly runs into difficulty when deployed across a global organisation with legacy infrastructure, regional variation, and real operational consequences. The people with direct experience of navigating that transition are scarce, and that scarcity is a practical constraint on how quickly organisations can move, regardless of appetite.

A separate concern, raised without a clear answer, is the pipeline question. If AI takes on the transactional, process-heavy work that has historically been where junior professionals develop their judgement and capability, the longer-term effect on the talent pool is uncertain. Some leaders were relatively relaxed about this, pointing to new roles that will require AI fluency rather than routine execution. Others were less sure. The transition period, before new role definitions settle, carries risk that is not yet well understood.

On the broader jobs question, one participant made a useful observation: the frictionless adoption assumption, that AI can be dropped into an organisation and simply produce results, does not survive contact with reality.

Meaningful technology adoption has always required sustained human effort to implement, maintain, and adapt over time. The fact that the major AI companies are now building their own professional services and forward-engineering functions is, in that reading, an acknowledgement that this deployment work is real and requires people to do it.

The CIO's Actual Problem

The suggestion that the CIO is becoming a de facto HR leader for AI agents was pushed back on fairly directly. The argument against was straightforward: agents are not employees, and the frameworks HR uses to manage people do not map cleanly onto the management of automated processes. Importing that language does not clarify anything useful.

What the discussion pointed to instead is less novel but perhaps more honest. The CIO's core responsibilities, process design, system interoperability, change management across complex environments, have not fundamentally changed. What has changed is the scale and speed of what is being managed, and the growing likelihood that external agents built by customers, partners, or others will begin interacting with an organisation's systems in ways that were not designed for and may not be immediately visible.

One concept that came up and deserves more attention is agent experience: treating the way autonomous agents interact with an organisation's systems as a design consideration, in the same way user experience and customer experience are. The argument for starting to think about this now is not that agent-to-agent interaction is already widespread in enterprise settings. It is not, yet. The argument is that the appetite for it is visible, the direction is reasonably clear, and the cost of preparing for it early is low relative to the cost of being unprepared when it arrives.

The Discipline Imperative

One line from the discussion cut through more than most: AI rewards discipline, not speed.

It is worth being specific about what that means in practice. Several leaders converged on a consistent set of criteria for evaluating agent use cases: does it address a genuine business need; can the value be measured in concrete terms; and is the right capability in place to build and scale it responsibly? One participant cited an example of a password-reset workflow rebuilt as an agentic process at a cost of six figures, a problem that a few lines of code would have solved. The failure there was not technical. It was the absence of basic commercial scrutiny before the work began.

The disruption concern that runs through these conversations is legitimate. Organisations with less legacy and lower operating costs are building multi-agent automated processes that will, in specific markets, compete with established enterprises. The clearest example offered was a US insurance carrier that built a fully agentic underwriting process for contracts below a defined value, not because the technology was interesting, but because it made a previously inaccessible market commercially viable. That is a business decision with an agentic solution, not a technology project in search of a justification.

For most enterprise leaders, the implication is not that they need to move faster. It is that the use cases worth pursuing are those where the commercial logic is clear before the technology conversation starts. Governance, data quality, and process ownership are not compliance overhead. They are the conditions under which scaling becomes possible. Building those foundations carefully is not the cautious option. It is the one that leaves room to accelerate later.

About CIONET

CIONET is the leading community of more than 10,000 digital leaders in 20+ countries across Europe, Asia, and the Americas. Through this global presence CIONET orchestrates peer-to-peer interactions focused on the most important business and technology issues of the day. CIONET members join over a thousand international and regional live and virtual events annually, ranging from roundtables, programs for peer-to-peer exchange of expertise, community networking events, to large international gatherings. Its members testify that CIONET is an impartial and value adding platform that helps them use the wisdom of the (IT) crowd, to acquire expertise, advance their professional development, analyse and solve IT issues, and accelerate beneficial outcomes within their organisation.