Beyond the AI Hype in HR: The Unspoken Reality

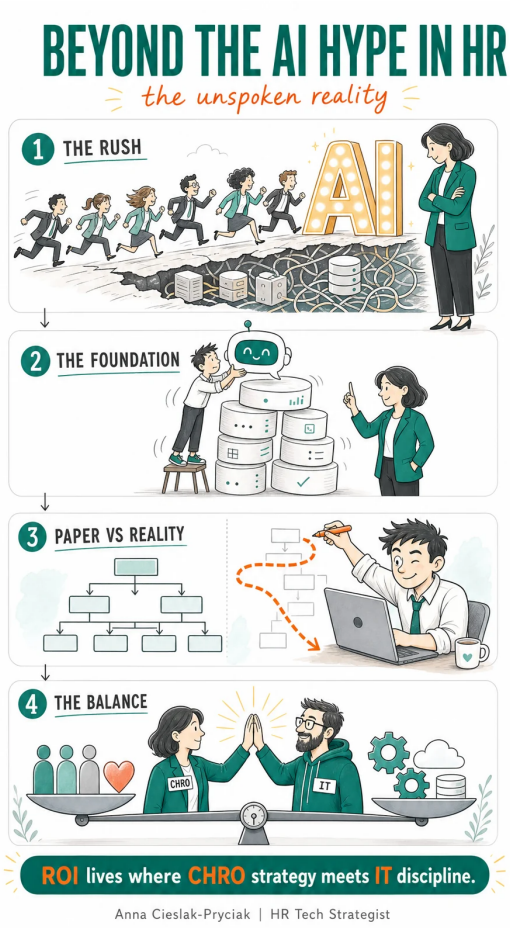
Anna Cieślak-Pryciak

HR Tech Strategist & Architect | HRTech Mood

The promise of Generative AI is no longer enough on its own. Business leaders need a clear view of what an investment should improve and how they will measure it. For CHROs and People Technology leaders, this means connecting business needs with the core IT architecture. Standalone tools can create fragmented data and additional silos when those connections are overlooked.

Sustainable value from AI in HR is co-architected at the intersection of HR strategy and IT infrastructure. Having spent many years architecting these ecosystems, I repeatedly see operational questions behind the technology questions: who owns the process, how the data flows, and what people actually need to do.

True, scalable ROI from AI in HR is co-architected at the intersection of HR strategy and IT infrastructure.



1. Moving from Isolated Features to True End-to-End Workflows

There is plenty of enthusiasm for implementing AI. In my experience, there is often less clarity around complete, end-to-end workflows and who owns them. A chatbot answering policy questions may be useful. The next question is how it fits into the wider service.

Two possible use cases illustrate the questions I would ask. These are examples to evaluate, not promised outcomes:

Retention analytics: could aggregated, appropriately governed signals help HR understand where support is needed? Any proposal needs scrutiny of data quality, fairness, privacy and how people will interpret the results. Its value must be tested against an agreed baseline.

Personalised onboarding: could role-relevant information and better sequencing help a new colleague become confident sooner? Define what readiness means, then measure the experience and outcome. A percentage improvement should come from evidence in that setting.

2. The Unspoken Foundations: Data Governance, Storage, and Security

You cannot build an intelligent enterprise on an unstable foundation. While the spotlight often shines on flashy AI interfaces, the real work happens in the background, focusing on how data is stored, structured, and protected. Before feeding massive amounts of employee data into any AI engine, People Tech and IT teams must establish absolute ownership over:

Data security and privacy: establish the purpose, permitted use, access controls and retention decisions before connecting employee data to AI. Involve the appropriate privacy, security and legal specialists.

Storage and Architecture: — Knowing exactly where the data resides, who has access to it, and how it syncs with your core HRIS ecosystem (such as Workday).

Weak data governance can undermine trust and introduce risk. In practice, I want to know who is accountable for the data, what quality checks are in place and what happens when something goes wrong.

3. The Paper vs. Reality Gap: Balancing Systems with Real Operations

Even when organizations know the importance of data and ownership, a massive problem remains continuously overlooked: the ownership of designing the actual processes. We

have all seen it — beautifully designed workflows mapped out on paper or in a slide deck. But in reality? People do things differently. Employees find workarounds, and managers bypass steps.

One of the most important aspects of my work has been balancing what can be configured in a system with how operations run in real life. When the two diverge, I look for the reason: a missing exception, an unclear owner, a training gap or a design that needs to change.

What I would prioritise

For me, sustainable value from AI depends on connecting process ownership, reliable data and architecture with a realistic understanding of human workflow. Start with an outcome, name the owner, test the complete process and learn from what actually happens.

Editorial noteThis is a practitioner perspective. The use cases are illustrative. No fixed ROI, retention saving or productivity improvement is claimed.