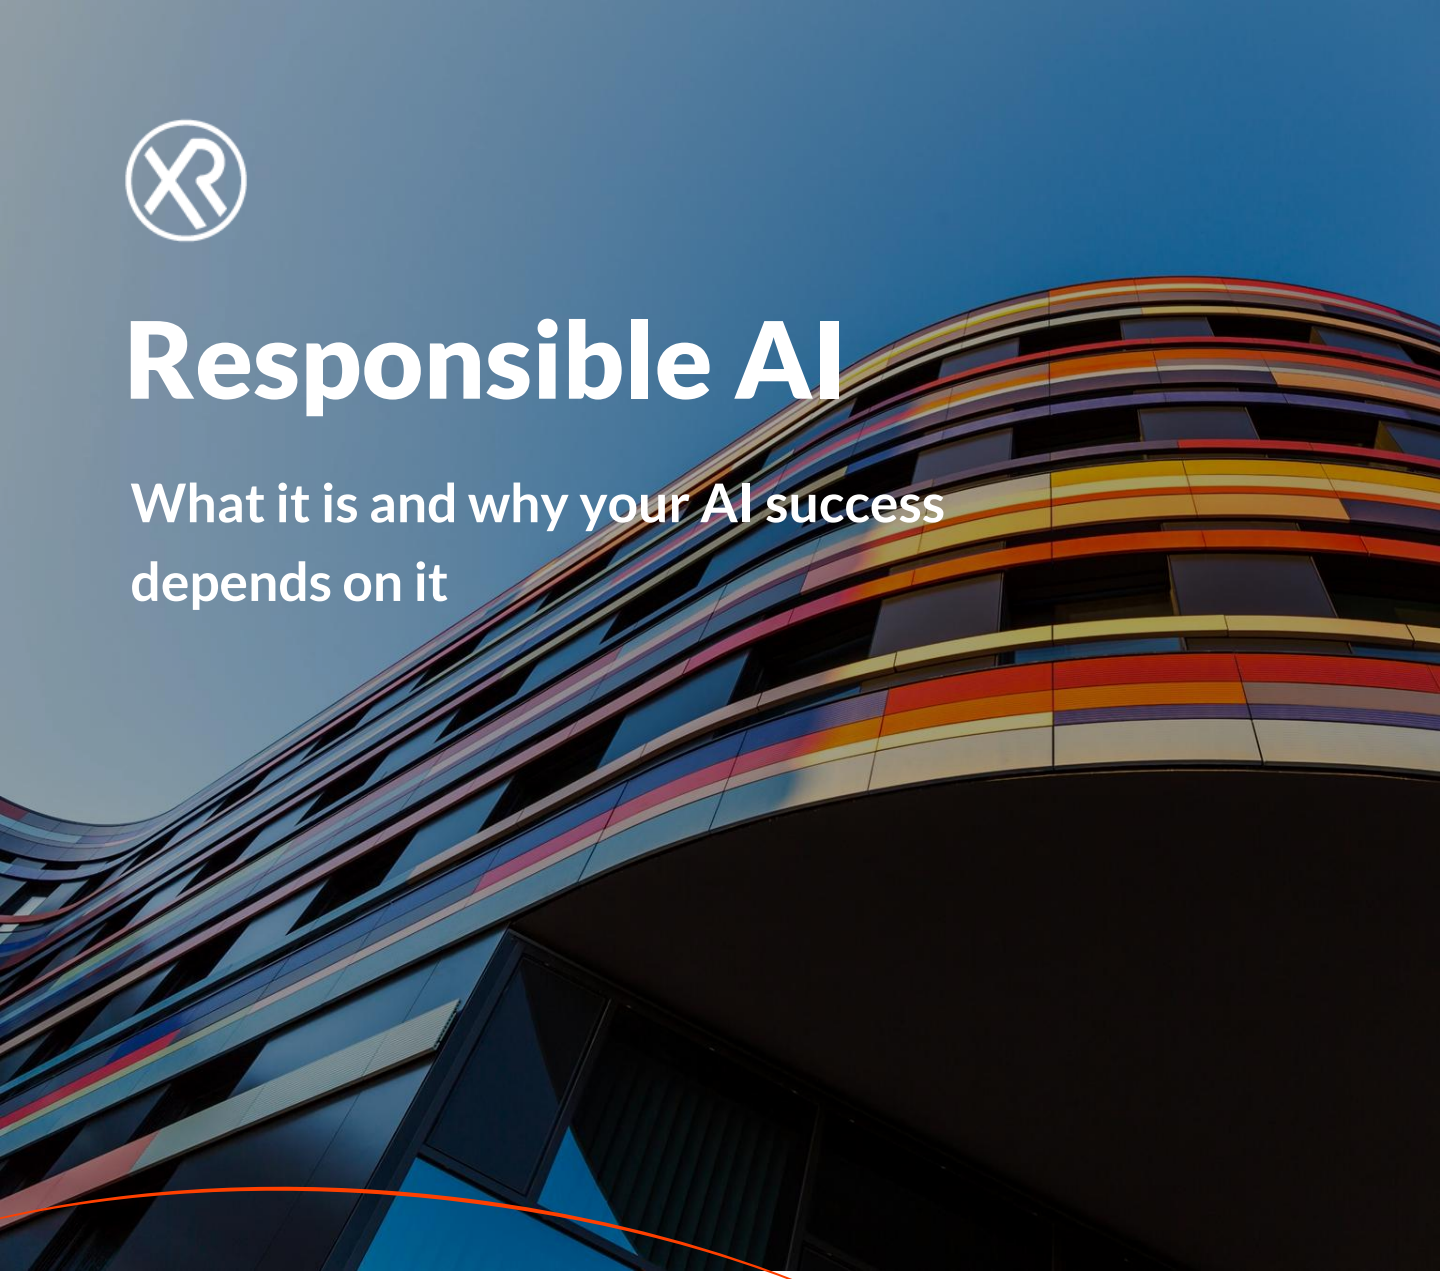

Responsible AI

What it is and why your AI success depends on it

xpertrule.com

Trust is the biggest barrier to successful AI investment.

The vast majority of AI projects are failing to deliver on budget or business goals, and it's not because the technology doesn't work.

It's because businesses can't trust the results.

Auditability, and the lack thereof, is the single biggest reason why 60% of AI initiatives are failing to deliver meaningfully on their business goals, according to McKinsey.

We can't scale what we can't trust — particularly in regulated sectors, where a single mistake could mean a catastrophic fine or a reputation ruined. The EU's AI Act, announced last year to kickstart a new regulatory era for AI, set out fines of €15 million or 3% of turnover for breaches in high risk cases. The stakes have never been higher, and our need for Responsible AI has never been greater.

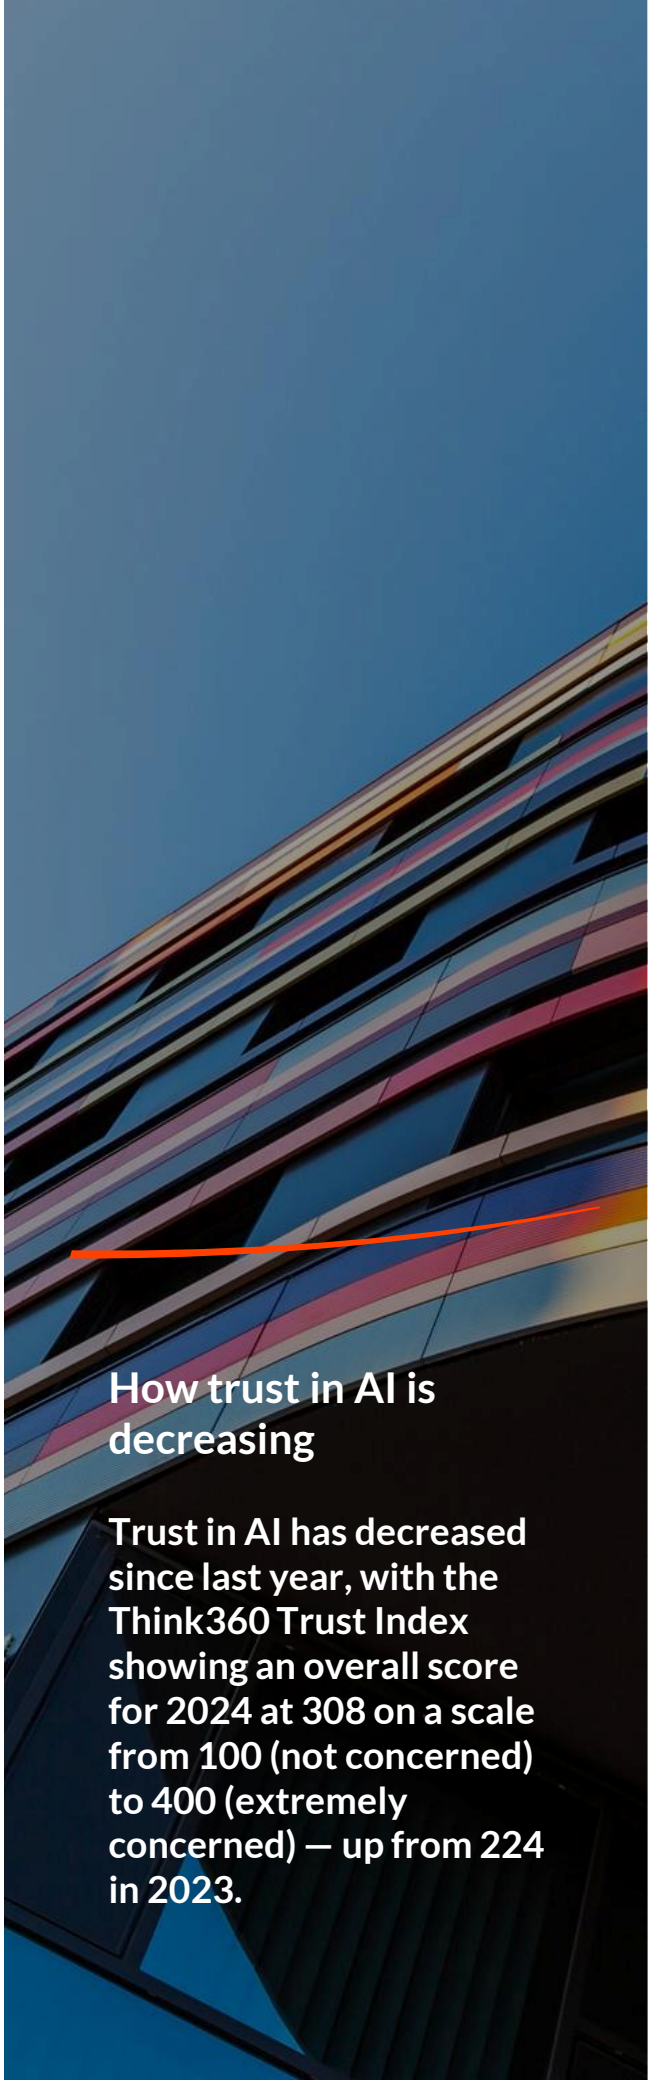

How trust in AI is decreasing

Trust in AI has decreased since last year, with the Think360 Trust Index showing an overall score for 2024 at 308 on a scale from 100 (not concerned) to 400 (extremely concerned) — up from 224 in 2023.

What constitutes Responsible AI?

Responsible AI is defined not by a single technology, but by a set of guiding principles for how AI solutions are designed and deployed: accuracy, consistency, accountability, auditability, compliance, safety and adaptability.

Ultimately, Responsible AI is AI you can trust with high stakes — or what we call mission-critical — decisions. It's AI you can readily audit, and deploy at scale without fear of hallucinations, bias, or incorrect decisions.

As AI proliferates, these errors — whether incorrect or biased decisions, data protection — are becoming increasingly common, with organisations including Microsoft, Air Canada, Tesla and Amazon all experiencing their own high profile incidents. The costs can be significant in terms of business continuity, brand damage, lawsuits and regulatory action.

Mission-critical decisions — the kind that require total trust and accountability — include the following use cases:

- Audit & due diligence

The XpertRule view on Responsible AI

"For complex processes and decision making, we must keep human experts in the loop. Not necessarily spending time on every simple task they may have traditionally carried out, but for critical decisions keeping human expertise and oversight in, is key. Additionally, we must prioritise transparency and auditability in the results AI applications are providing us to protect against bias."

Aden Hopkins, XpertRule CEO

- Structured inspection & reporting
- Fraud detection
- Insurance claims processing
- Resource planning and optimisation
- Performance optimisation
- Risk assessment
- Complex & regulated diagnostic and troubleshooting

It's in these consistently high-risk decisions where AI's impact can be most transformative, with the average S&P company losing $250 million annually due to poor decision-making affecting business continuity.

And it's in this space, with the stakes at their highest, where Responsible AI is needed most. In fact, demand for "explainable, auditable AI" has grown 6x year-over-year in RFPs across heavily regulated industries like Pharma and FinServ (IDC Decision Tech Trends, Q4 2023).

In response to this market gap, "built-in transparency" is now one of the top three criteria in enterprise AI procurement, according to Forrester's Responsible AI Buyer Survey in 2023. The problem enterprises face is that built-transparency is not what most AI vendors are offering.

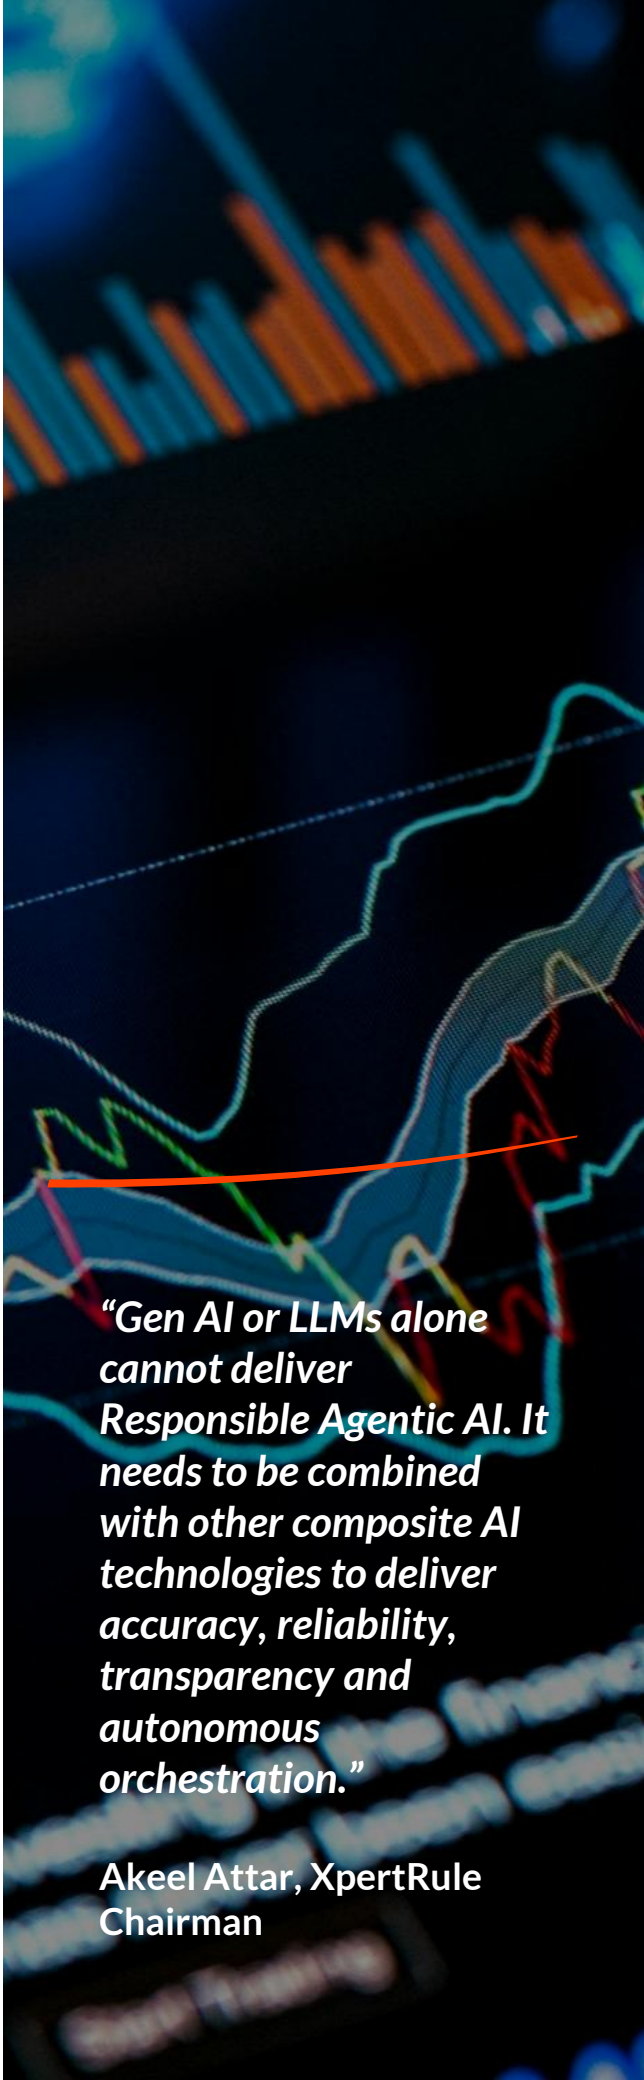

"Gen AI or LLMs alone cannot deliver Responsible Agentic AI. It needs to be combined with other composite AI technologies to deliver accuracy, reliability, transparency and autonomous orchestration."

Akeel Attar, XpertRule Chairman

Treating transparency as an afterthought is a recipe for disaster

Retrospective transparency — that is, tools that explain their AI-powered decisions after the decision has been made — has become the norm.

Most decision intelligence platforms rely on retrospective explanations, which can be biased or incomplete, says the Harvard Business Review. But justifying outcomes after the fact does not equate to actual transparency

That's not just how the technologists see it. Regulators are increasingly skeptical of black-box AI models even when paired with explanation tools, according to Stanford HAI's 2023 EU AI Act Analysis.

When most vendors say they validate the output of Gen AI, they're in fact using it to make a high-level decision or classification, then relying on retroactive, handcrafted rules or other Gen AI modules to attempt to validate the output.

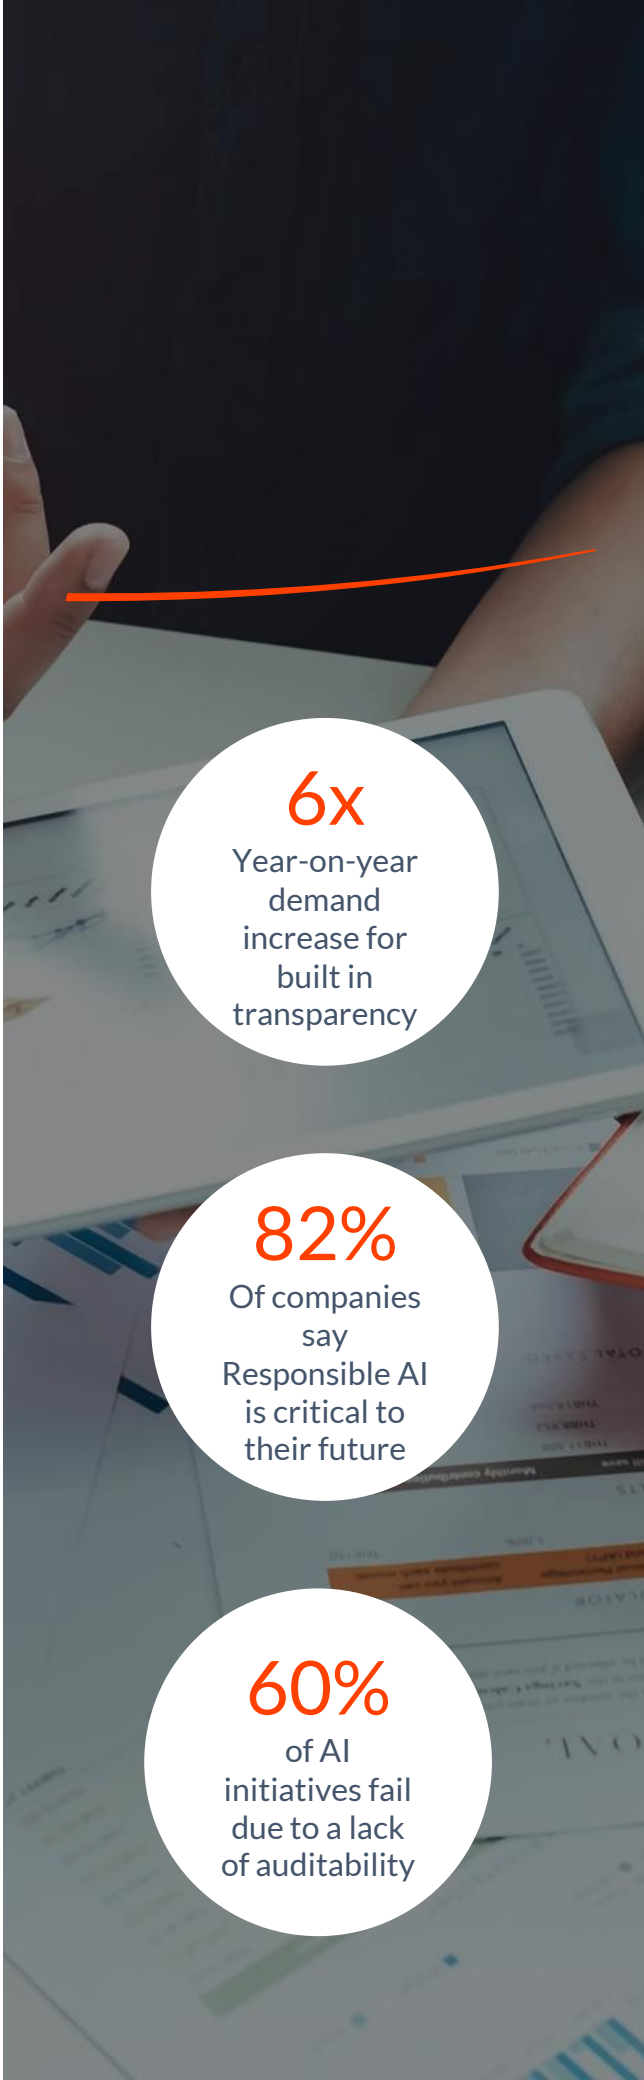

Most AI agents or solutions being deployed today are powered by Large Language Models (LLMs). This means that the AI agents being deployed are failing the Responsible AI test – specifically for the following reasons:

- LLMs are non-deterministic systems—essentially black boxes that make it difficult for human users and engineers to determine when an agent has achieved its goal and whether it performed as expected.
- LLMs still face issues with hallucination, where the model output is fabricated and not consistent with the provided context or the world knowledge of the model. In agentic systems where LLMs are creating / orchestrating the sub-tasks to achieve their goals, underlying hallucinations can lead to compounding mistakes in agent behaviour as they are carried through the workflow.
- LLM-powered agents have limited planning, reasoning, and real-time learning capabilities due to their inability to build an accurate model of the real world and of the enterprise. Furthermore they lack common sense, temporal, and causality reasoning.

The resulting challenge for a lot of businesses is that they don't realise their AI solution isn't responsible until it's too late — until the fine has been issued, the headlines are written and the reputational or operational (or both) damage is done.

60% of AI failures are tied to a lack of auditability and poor alignment with governance/compliance frameworks.

McKinsey Tech Trends Report, 2023

We design AI solutions with auditability at their heart

Compliance requires real-time traceability, not just rationalisation after decisions are made.

XpertRule have developed a Composite AI Decision Intelligence platform that allows the development of Responsible AI Agents that overcome the above limitations of LLM powered agents and deliver the enormous potential of agentic AI in transforming work, services and enterprise applications.

With all of our AI solutions, expert validation and approval are baked into design. We use decision intelligence power prompting, meaning that the design and sequencing of prompts is intelligently guided by an inference engine. This ensures Gen AI is always invoked, with full context and human oversight, at the lowest possible micro level — before any damage has been done.

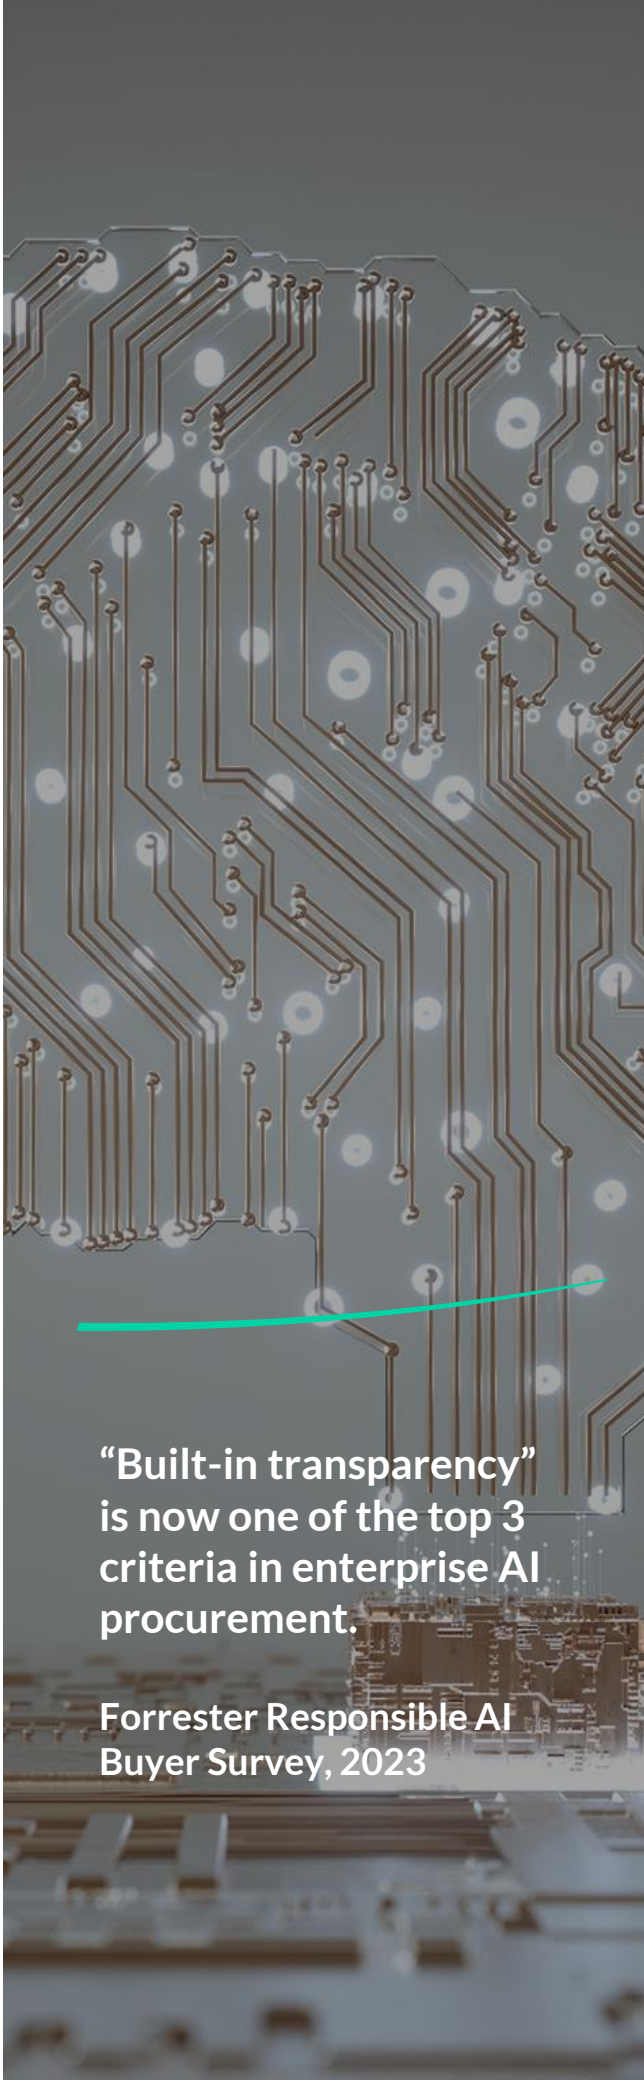

“Built-in transparency” is now one of the top 3 criteria in enterprise AI procurement.

Forrester Responsible AI Buyer Survey, 2023

Responsible AI is technology that puts human expertise first

Responsible AI is not just about guarding businesses against the consequences of biased or incorrect decisions. It's about deploying AI with human interests and ethics at heart.

Our solutions are designed to protect, maximise and champion human expertise. Having witnessed countless boom-and-bust cycles, we've learned that a blanket tech-first approach to AI implementation crowds out human potential, diminishes opportunities for career progression and actually decreases the quality of business outcomes.

Instead we take a decision-centric approach to automation – supporting the capture and maintenance of large scale, complex knowledge bases that help people make better decisions and drive better business outcomes.

Regulators are increasingly skeptical of black-box AI models even when paired with explanation tools.

EU AI Act Analysis, Stanford HAI, 2023

This is because after 40 years in the business, we know better than anyone the value of human expertise. The knowledge of an organisation lies in the brains of its best people – but when these people leave the workforce, their knowledge goes with them. Our solutions capture specialist decision making knowledge, use it to drive consistently better outcomes, and keep it inside the business, forever.

Keeping human expertise in the loop is an essential component of Responsible AI — not just for the accuracy and quality of outcomes, but also for the human experience of working alongside AI. A lack of trust leads to the overreliance on, and even exploitation of, the human in the loop. If AI can't be trusted, it falls to human experts to laboriously validate every automated action with fact-checking — and as a result, AI fails to deliver on its promise of fairer, better jobs.

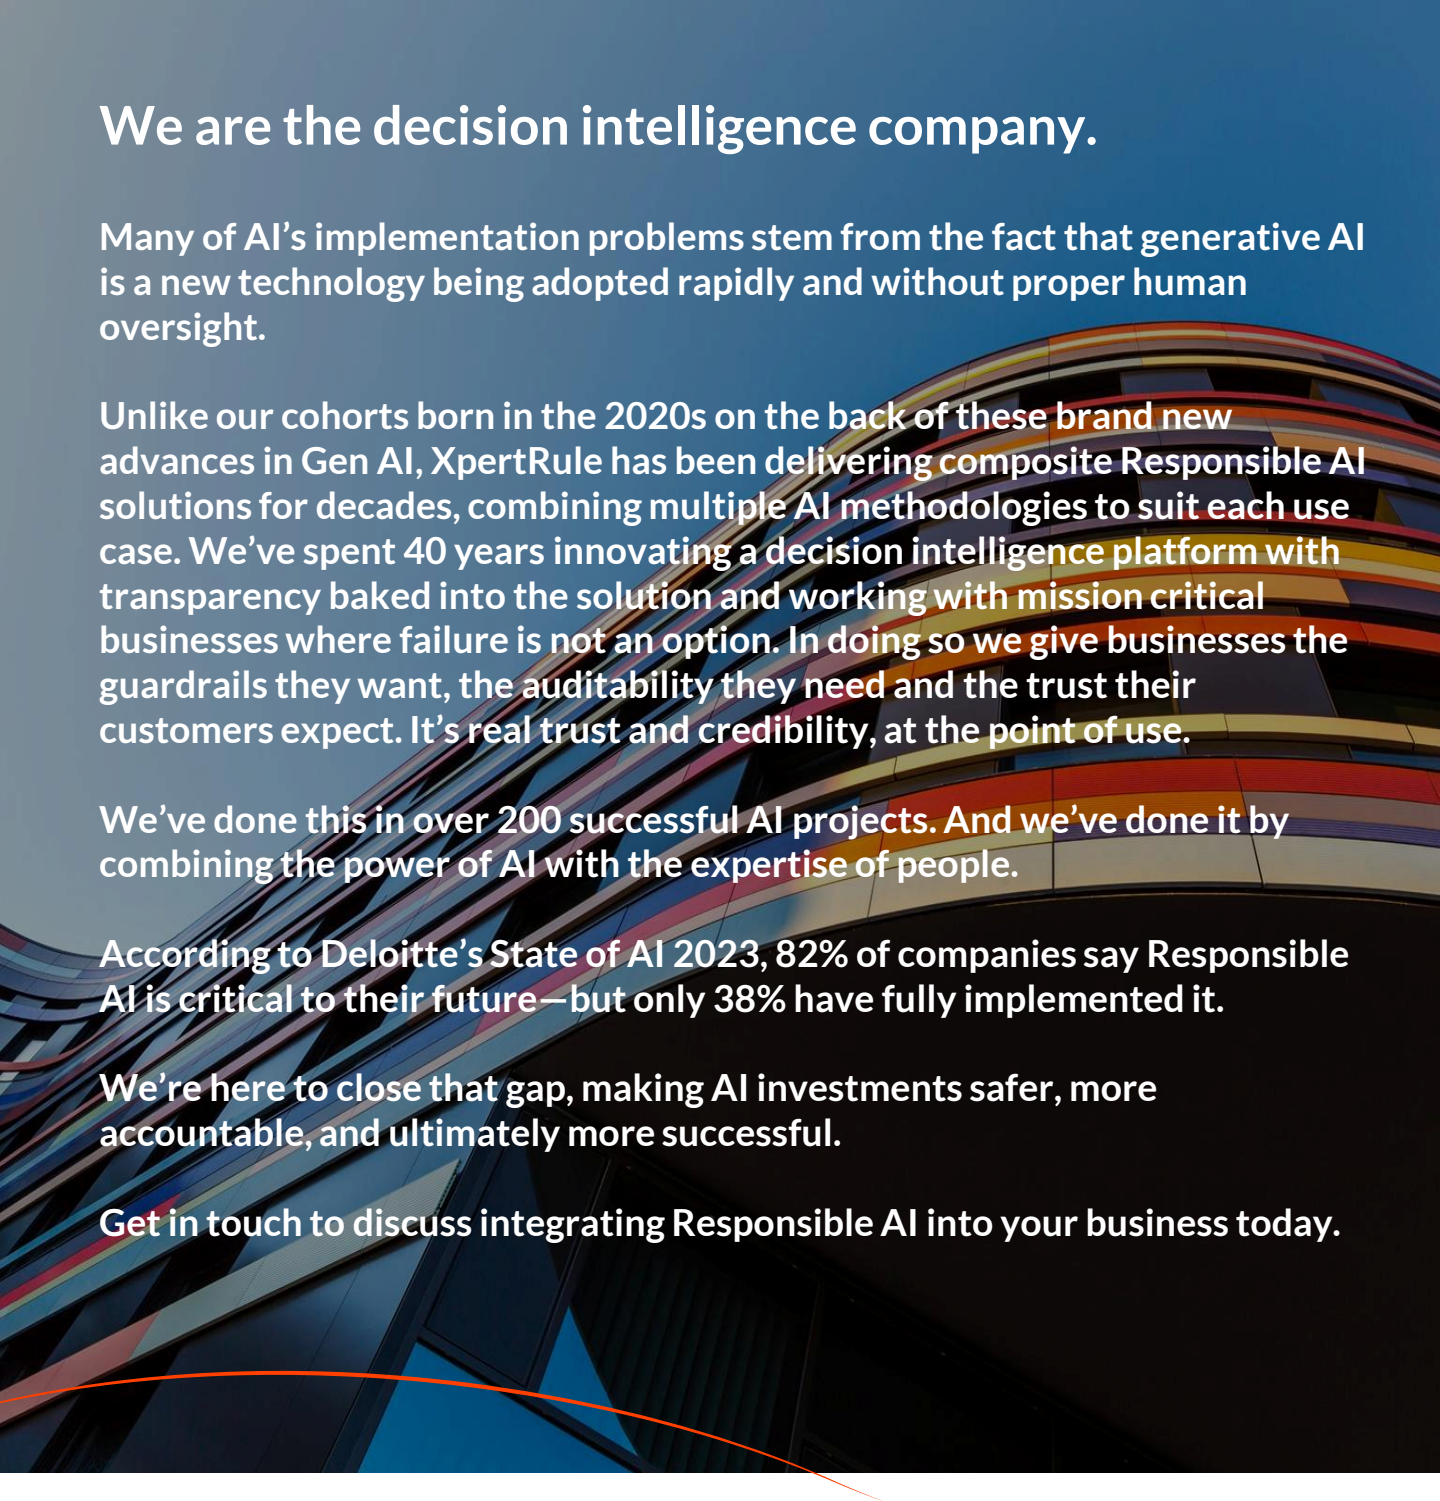

We are the decision intelligence company.

Many of AI's implementation problems stem from the fact that generative AI is a new technology being adopted rapidly and without proper human oversight.

Unlike our cohorts born in the 2020s on the back of these brand new advances in Gen AI, XpertRule has been delivering composite Responsible AI solutions for decades, combining multiple AI methodologies to suit each use case. We've spent 40 years innovating a decision intelligence platform with transparency baked into the solution and working with mission critical businesses where failure is not an option. In doing so we give businesses the guardrails they want, the auditability they need and the trust their customers expect. It's real trust and credibility, at the point of use.

We've done this in over 200 successful AI projects. And we've done it by combining the power of AI with the expertise of people.

According to Deloitte's State of AI 2023, 82% of companies say Responsible AI is critical to their future—but only 38% have fully implemented it.

We're here to close that gap, making AI investments safer, more accountable, and ultimately more successful.

Get in touch to discuss integrating Responsible AI into your business today.

Get in touch.

xpertrule.com