veeam

Data and AI Trust Maturity Model

Finding Your Path to Safe AI At Scale

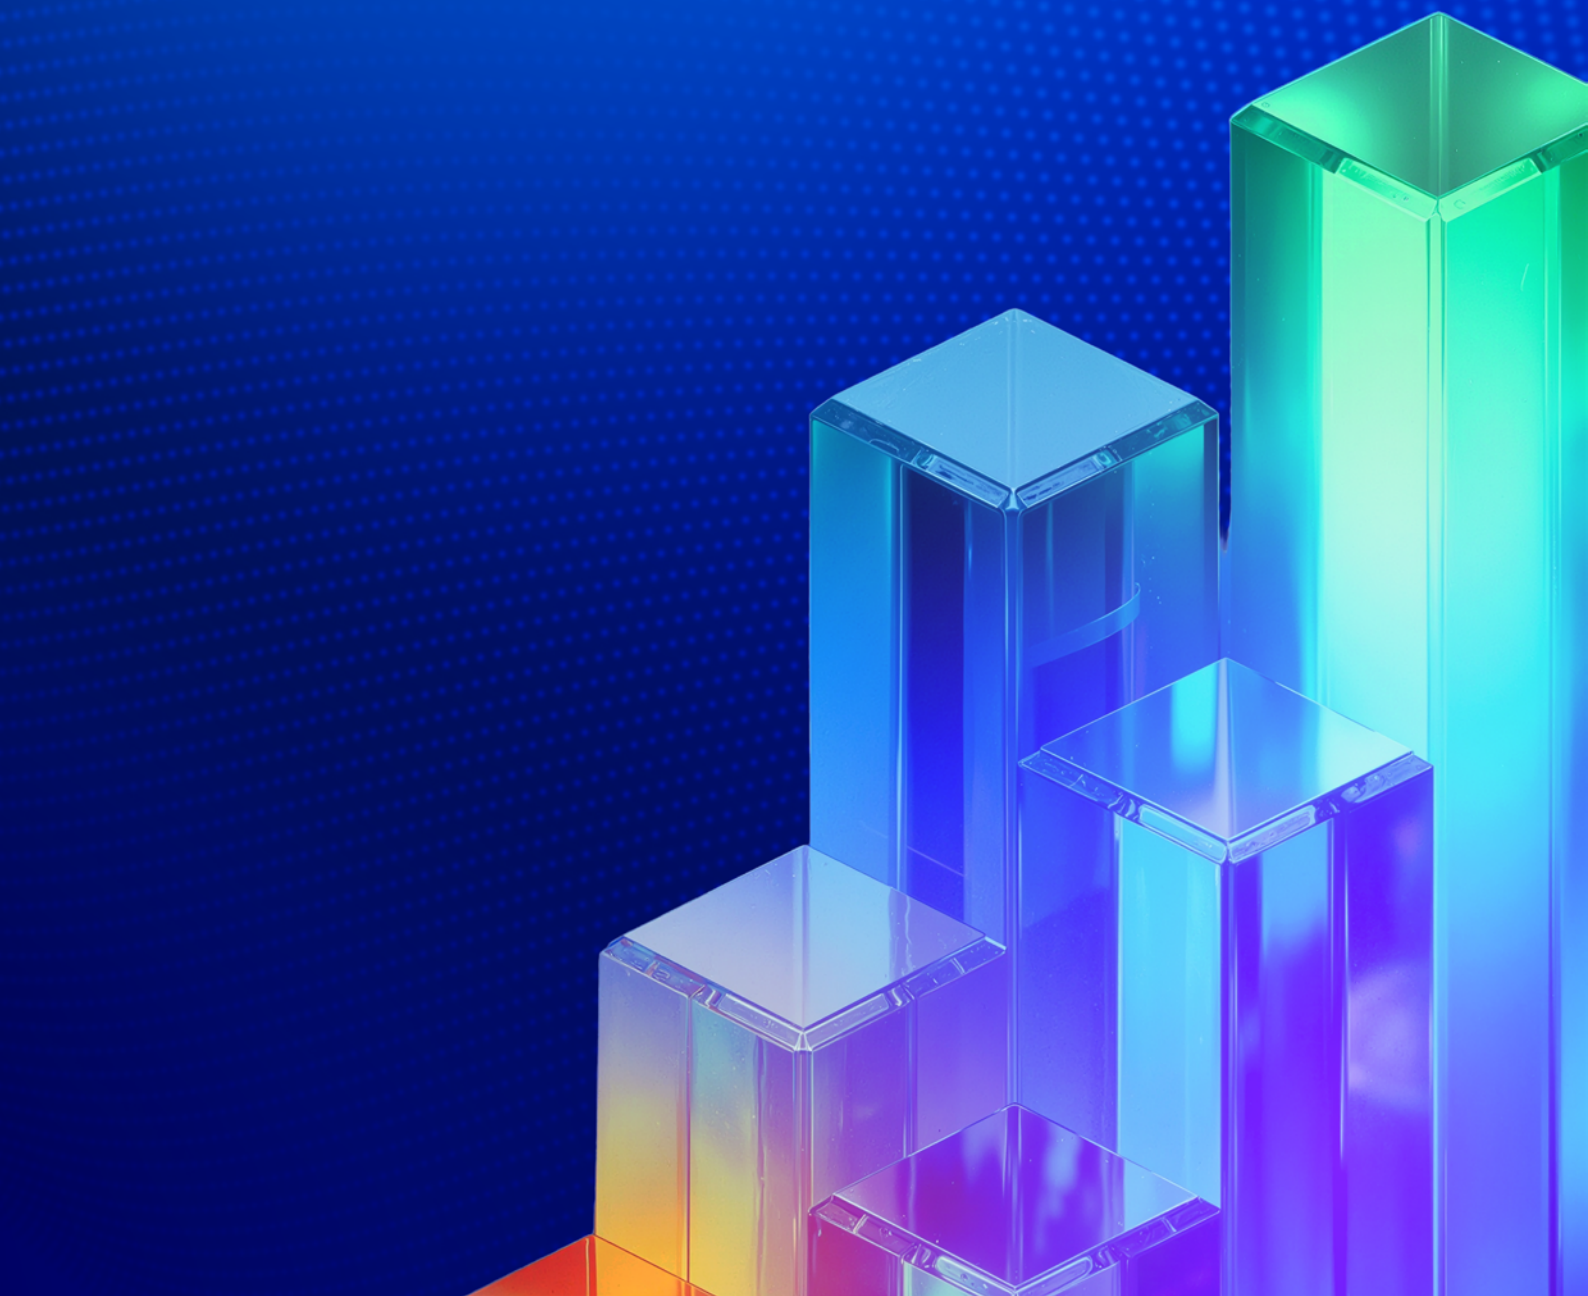

Finding Your Path to Safe AI At Scale

This paper draws from an independent research study of 300 technology and business leaders across financial services, healthcare, government, manufacturing, and technology sectors. Respondents include senior decision-makers with direct accountability for data, security, risk, and technology strategy. The findings specifically reflect responses from C-suite leaders across various industries, as they are the ones closest to the decisions that determine whether AI scales safely or stalls.

There is a significant gap between deploying AI and being ready for what that deployment demands. Across the globe, companies rank in vastly different ways in terms of their preparedness for AI deployment and all it entails.

Our Data & AI Trust Maturity Model research report has split organizations across five different levels, ranking their deployment, implementation, and preparation. Our research found that 95% of organizations have deployed AI in their environment. However, deploying AI safely and effectively is proving more difficult, because the operational foundations required to trust AI at scale take longer to build than the technology takes to adopt.

AI is Scaling Fast — But So Are the Setbacks

AI adoption among these respondents' organizations is well underway. Nearly 7 in 10 say AI is either already embedded across multiple business functions or is central to their operations and competitive strategy. Yet the same population reports a pattern of setbacks that is difficult to reconcile with that level of maturity.

Around **52%** of respondents say at least one AI initiative was scaled back in scope over the past 18 months. 4 in 10 report initiatives delayed beyond their original timelines. More than 1 in 4 **(28%)** report an initiative discontinued entirely, while about 1 in 3 say their initiatives have proceeded as planned.

These setbacks are occurring in organizations where AI is already in production, which means they have less to do with early-stage adoption and more to do with what adoption requires once it reaches scale.

Speed Exposes the Gaps

The barriers leaders identify help explain why. The most frequently cited constraint is talent gaps in AI and machine learning expertise **(43%)**, followed by difficulty integrating AI into existing workflows and systems **(33%)**. Regulatory uncertainty **(25%)**, data quality limitations **(20%)**, and explainability concerns **(19%)** add further pressure.

These challenges are not independent of one another. Rather, they compound in organizations that move quickly into AI deployment before governance, accountability structures, and data infrastructure are ready to support it. The barrier list, in that sense, is less a set of discrete problems than a description of what gets left behind when speed takes priority.

Notably, few executives cite the technology itself as the limiting factor, reinforcing that AI maturity is now constrained more by operational readiness than innovation velocity.

Where AI Hits Limits:
Skills, Systems, and Standards

Executive confidence, on the surface, tells a more optimistic story. Some **80%** of executives say they are confident in their organization's ability to scale AI safely and responsibly over the next two years. The research, however, reveals a critical distinction: nearly half of all respondents acknowledge that their confidence is driven more by intuition than by evidence they could readily demonstrate to an external audience. Among executives whose confidence is grounded in a definitive plan, three factors carry almost equal weight:

55%

A formal governance framework

55%

A visible leadership commitment

53%

A dedicated AI risk function

However, the numbers start to drop when it comes to external validation. Only 31% cite having passed a regulatory review or audit involving AI, and just **20%** reference benchmarking against industry peers. For most organizations, confidence is internally generated rather than externally tested. It reflects the frameworks leaders have designed, the teams they have built, and the priorities they have set — all of which must be pressure-tested externally to prove they can hold up in an emergency.

This distinction matters as expectations for AI governance evolve. Regulators and boards are increasingly less concerned with whether governance exists in principle and more focused on whether it can be demonstrated on demand.

Nearly 9 in 10 executives report that formal AI governance policies exist in some form. Yet only about 1 in 3 say they could produce comprehensive audit evidence immediately if asked. Most acknowledge that the evidence exists, but it would require substantial effort to compile. Written policies and auditable governance are distinct capabilities and the gap between them represents a meaningful source of exposure.

Accountability structures reinforce this pattern. Most executives place primary responsibility for AI governance with the CTO or CDO, while a significant share rely on cross functional committees to coordinate oversight. **7%** have a dedicated AI governance leader.

While each of these models can be effective, they depend on clearly-defined ownership and accountability, two conditions that become harder to sustain as AI expands across functions. Without explicit roles and escalation paths, policy coverage does not reliably translate into operational control, particularly when something goes wrong at scale.

Research suggests an environment in which AI adoption is widespread and leadership optimism is high. However, the operational foundations needed to sustain adoption and enthusiasm are still taking shape in most organizations. The companies navigating this transition most effectively share a defining characteristic: confidence grounded in proof. In particular, organizations ranked in the upper levels of this report (levels four and five) tend to have a practical governance program that has been implemented and tested under scrutiny, not merely designed.

AI at Scale Favors Preparation, Not Speed

The Data & AI Trust Maturity Model provides leaders with a structured, benchmarked way to close the gap between perception and reality. It assesses how consistently an organization governs, secures, recovers, and enables data and AI across the enterprise, measuring operational reality rather than stated intent. Built on research and leading practices, the model measures five maturity levels, from ad hoc to industry leading. It turns readiness from an assumption into a measurable capability.

For leaders who question whether their confidence is ahead of their evidence, it also provides the external reference point the research suggests most organizations are missing.

The organizations most likely to sustain AI at scale are not necessarily the fastest adopters. Instead, they are the ones that invest early in governance, accountability, and data foundations that allow AI to scale without introducing risk that outpaces their ability to manage it.

This executive summary highlights data from the US region. The full Data & AI Trust Maturity Model research report expands on these findings across more regions and industries, and examines how leaders approach governance, resilience, and AI readiness. This report gives leaders a way to assess where they stand against that standard and provides a clear path to meeting it.

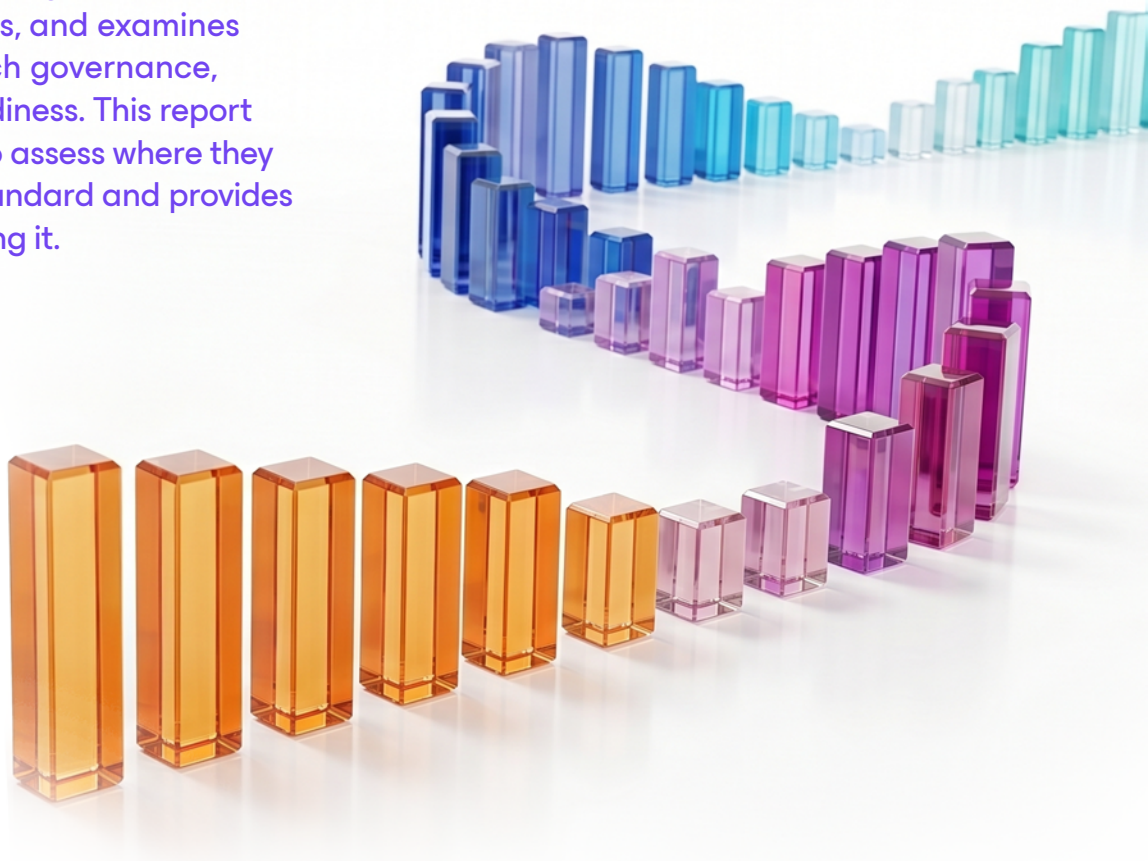

Methods & Scope

This executive summary is based on Veeam-commissioned research conducted with 300 technology and business leaders across financial services, healthcare, government, manufacturing, and technology. Respondents included senior decision-makers with responsibility for data, security, risk, and technology strategy. The findings referenced throughout this paper focus on C-suite respondents within the broader study population.

The research examined how organizations are adopting, governing, and scaling AI in practice, with a focus on operational readiness, governance structures, and the conditions that shape leadership confidence. Findings reflect reported behaviors, outcomes, and decision-making conditions rather than stated future intent. Benchmark results are drawn from the full study and represent actual responses from participating leaders, used to identify patterns common across organizations as AI moves from initial deployment into broader business operations.

About Veeam Software

Veeam is the Data and AI Trust Company, specializing in helping organizations ensure their data and AI are fully understood, secured, and resilient to enable the acceleration of safe AI at scale. As the market leader in both data resilience and data security posture management, Veeam is built for the convergence of identity, data, security, and AI risk.

Veeam delivers deep contextual intelligence across every data asset, identity, and AI model. The company governs access for both humans and AI agents, automates privacy, compliance, and remediation processes, and protects and recovers organizations from modern threats — including ransomware, disasters, AI errors, and ensuring the restoration of clean, trusted data. Veeam empowers organizations to move beyond simply protecting data, enabling them to activate and unlock its full potential.

Headquartered in Seattle with offices in more than 30 countries, Veeam protects over 550,000 customers worldwide, including 82% of the Fortune 500, who trust Veeam to keep their businesses running.

Learn more at www.veeam.com or follow Veeam on LinkedIn @veeam-software and X @veeam.