altman solon

SOFTWARE AI RESILIENCY:

A Strategic Framework for Cybersecurity Investors

AI is shaping a new era in the cybersecurity landscape by accelerating automation across parts of the security workflow, shifting where meaningful differentiation lies in the security stack. AI enables feature-level capabilities such as scanning, triage, and rule generation to become easier to replicate. At the same time, value is increasingly concentrated in proprietary data signals, deep workflow integration, and the ability to navigate trust, regulation, and operational risk. Here, we apply Altman Solon's AI Resilience Diagnostic framework to cybersecurity software, drawing on proprietary research with over 50 Chief Information Security Officers (CISOs) and cybersecurity practitioners in the U.S., Europe, and the U.K. to identify which categories are more exposed to disruption and where sustainable value remains, offering opportunity for outsized performance.

Investor interest in AI-enabled cybersecurity remains high, but conviction is uneven as AI reduces replication costs and blurs moats across the security stack. In this environment, AI resiliency is quickly becoming a decisive determinant of investment conviction and confidence in future value creation. Many investors recognize the opportunity but struggle to distinguish between businesses that AI will strengthen and those that may be weakened.

In our earlier paper, we introduced a moat-first framework for assessing software AI resiliency. Here, we apply the same two-tier framework to cybersecurity software. In this market, **automation and code generation alone will not determine value. Rather, access to proprietary data, highly regulated workflows, and durable technical and non-technical moats will be critical to sustainable value**.

Where is AI headed in cybersecurity?

AI is moving deeper into cybersecurity workflows, from assistive tools and task-level automation toward embedded decision support across security operations, code security, and remediation. A new wave of AI-native security companies is beginning to automate parts of the security workflow itself. A growing set of tools now embed large language models directly within developer and security workflows to analyze codebases, triage alerts, or coordinate remediation steps. Examples include Claude Code by Anthropic, autonomous SOC agents like 7AI, and AI-enabled SecOps platforms such as Abstract Security. These tools do not eliminate cybersecurity categories, but they do accelerate automation of repeatable tasks. Platforms such as SIEM and XDR remain necessary for integrating signals, orchestrating response actions, and providing accountability across complex enterprise environments.

AI is not only compressing parts of the security stack, but it is also increasing the scale and speed of the problem. Software development cycles are accelerating, attacker automation is improving, and security telemetry is expanding across cloud, identity, endpoints, and SaaS systems. As signal volumes grow, organizations will still rely on platforms that can aggregate signals, coordinate response workflows, and enforce governance across those environments, even as AI automates many of the underlying tasks.

AI is also accelerating convergence across cybersecurity categories as standalone products increasingly merge into broader platforms to support automated detection, triage, remediation, and governance workflows. Security operations tools integrate threat detection, investigation, and response. Cloud security platforms combine posture management, workload protection, and application security, and identity signals increasingly feed into broader security decision-making frameworks. AI reshapes not only individual product capabilities but also the way categories are defined.

Categories mapped to the NIST cybersecurity lifecycle

High-level, not exhaustive, categories may span multiple lifecycle functions

- Less vulnerable
- Somewhat vulnerable, use case specific
- Most vulnerable

	IDENTIFY	PROTECT	DETECT & RESPOND	RECOVER
PERSPECTIVE	Remains comparatively resilient due to deterministic controls, governance requirements, and regulation	Undergoing AI-driven transformation, with vulnerability highest where workflows are more easily automated or commoditized	SecOps capabilities are being transformed by AI-driven automation across detection, investigation and response workflows. MDRs without proprietary data or vertical focus risk displacement by automation-first competitors and 'platformization' trends	Remains comparatively defensible due to complexity and the mission-critical nature of recovery capabilities
SELECTED SUB-CATEGORIES	**VERIFICATION & ACCESS** IAM, PAM, authentication, and IGA platforms, and services	**CLOUD SECURITY** Asset discovery, ASM, and VM	**DETECTION PLATFORMS (SIEM/XDR)** Security analytics platforms aggregating telemetry and identifying threats	**BACKUP AND CYBER RESILIENCE** Backup platforms, disaster recovery orchestration and resilience tooling, breach remediation
		DATA LOSS PREVENTION Sensitive data management and governance		
		NETWORK DEFENSE Firewalls, gasteways, network security controls	**MANAGED DETECTION & RESPONSE / SOC** Managed monitoring, investigation, and threat response services	
	EXPOSURE MANAGEMENT Asset discovery, attack surface monitoring, and vulnerability prioritization platforms	**ENDPOINT PROTECTION** EDR, EPP		
		APPLICATION SECURITY / DEVSECOPS SAST, DAST, code analysis, dev tooling	**INCIDENT RESPONSE & FORENSICS** Investigation, containment, and remediation of security breaches	
		EMPLOYEE TRAINING SAT platforms and behavioral training programs		

Source: Altman Solon

Despite rapid progress, a large proportion of cybersecurity use cases still require human supervision at critical points in the workflow, and many security environments remain difficult to automate end-to-end. For example, in regulated environments, opaque decision-making can create legal, compliance, and forensic risk. In practice, this means that **cybersecurity is unlikely to become fully autonomous in the near future**. Instead, the market is shifting toward hybrid models in which AI augments analysts, accelerates workflows, and compresses some forms of product differentiation faster than others.

For investors, the implication is clear: **AI will not affect all cybersecurity categories equally**. The key question is not whether AI will reshape the market, but where it will erode standalone value, where it will reinforce existing moats, and which companies are poised to benefit.

Cybersecurity investors should evaluate AI resiliency at the category and the company level

To identify sustainable moats in cybersecurity software, investors must evaluate both the category and the company. If a category is more exposed to AI automation, the specific company must have stronger technical or non-technical moats to compensate.

AI-resilient cybersecurity categories often share five characteristics:

1. **Variable, complex processes**:
   Situations requiring analyst judgment and adversarial reasoning.
2. **Qualitative, unstructured inputs**:
   Fragmented or regulated data that AI struggles to generalize.
3. **Limited training data**:
   Niche environments where "big data" models lack sufficient context.
4. **Intolerance to probabilistic output**:
   Missions where the "cost of error" is too high for a model that might hallucinate.
5. **Unclear/variable success criteria**:
   Environments where the "right" answer shifts as adversaries adapt.

These characteristics translate directly into observable differences in exposure to disruption across the cybersecurity market.

Our survey of over 50 CISOs and senior security operations experts found that **security operations & detection and application & code security are the most exposed**, with 48% of respondents identifying them as vulnerable to AI-driven disruption. Both categories are heavily weighted toward repeatable, rules-based workflows, exactly the conditions where AI automation erodes standalone value most quickly. Respondents cited two reinforcing dynamics: AI-driven automation compressing the value of signature-based tools, and platform bundling pulling adjacent capabilities into broader ecosystems. Managed detection and response (MDR) services sit somewhere in the middle along this spectrum. While AI can automate portions of investigation and alert triage, the service model still relies on human expertise, customer-specific tuning, and trusted response workflows, making disruption exposure more use-case dependent.

Other categories appear more resilient, such as endpoint protection (27%) and network security (31%), which benefit from deep hardware integration, OS-level telemetry, and embedded data advantages that are difficult for new entrants or AI alone to replicate. **Categories anchored in trust, deterministic logic, and regulatory accountability show the most resilience**: identity & access management (IAM) remains one of the strongest areas, and backup and recovery benefit from their non-negotiable role in business continuity.

To invest in sustainable moats, ask:

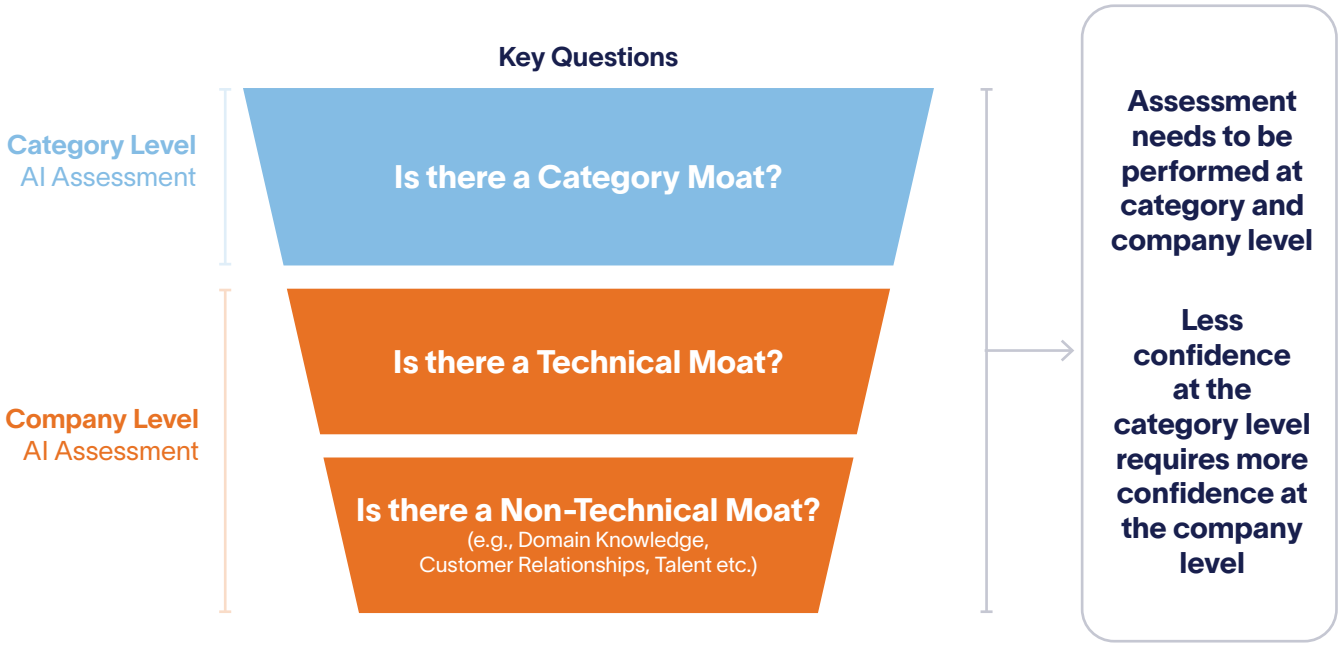

Source: Altman Solon

When asked what drives durable standalone value, respondents converged on three core sources of defensibility:

1. **Access to proprietary data**, specifically, large-scale telemetry aggregated across customer environments rather than signals owned by any single organization (63%).
2. Use cases where a probabilistic AI answer is legally or operationally insufficient, such as **regulated compliance workflows** (60%).
3. Systems that require a **deep understanding of a specific organization's intent and risk appetite** (52%).

Taken together, we see a clear pattern: the categories most exposed to disruption are those where AI can quickly replicate functionality, while the most resilient are those where data, trust, and context create barriers to automation.

Applying these characteristics across cybersecurity categories reveals meaningful variation in exposure to AI-driven disruption. **Categories built around repeatable, rules-based workflows tend to face greater automation pressure, while those anchored in trust, regulatory requirements, or deep system integration tend to be more resilient**. In practice, most categories fall somewhere between these extremes, with company-specific differentiation often proving decisive.

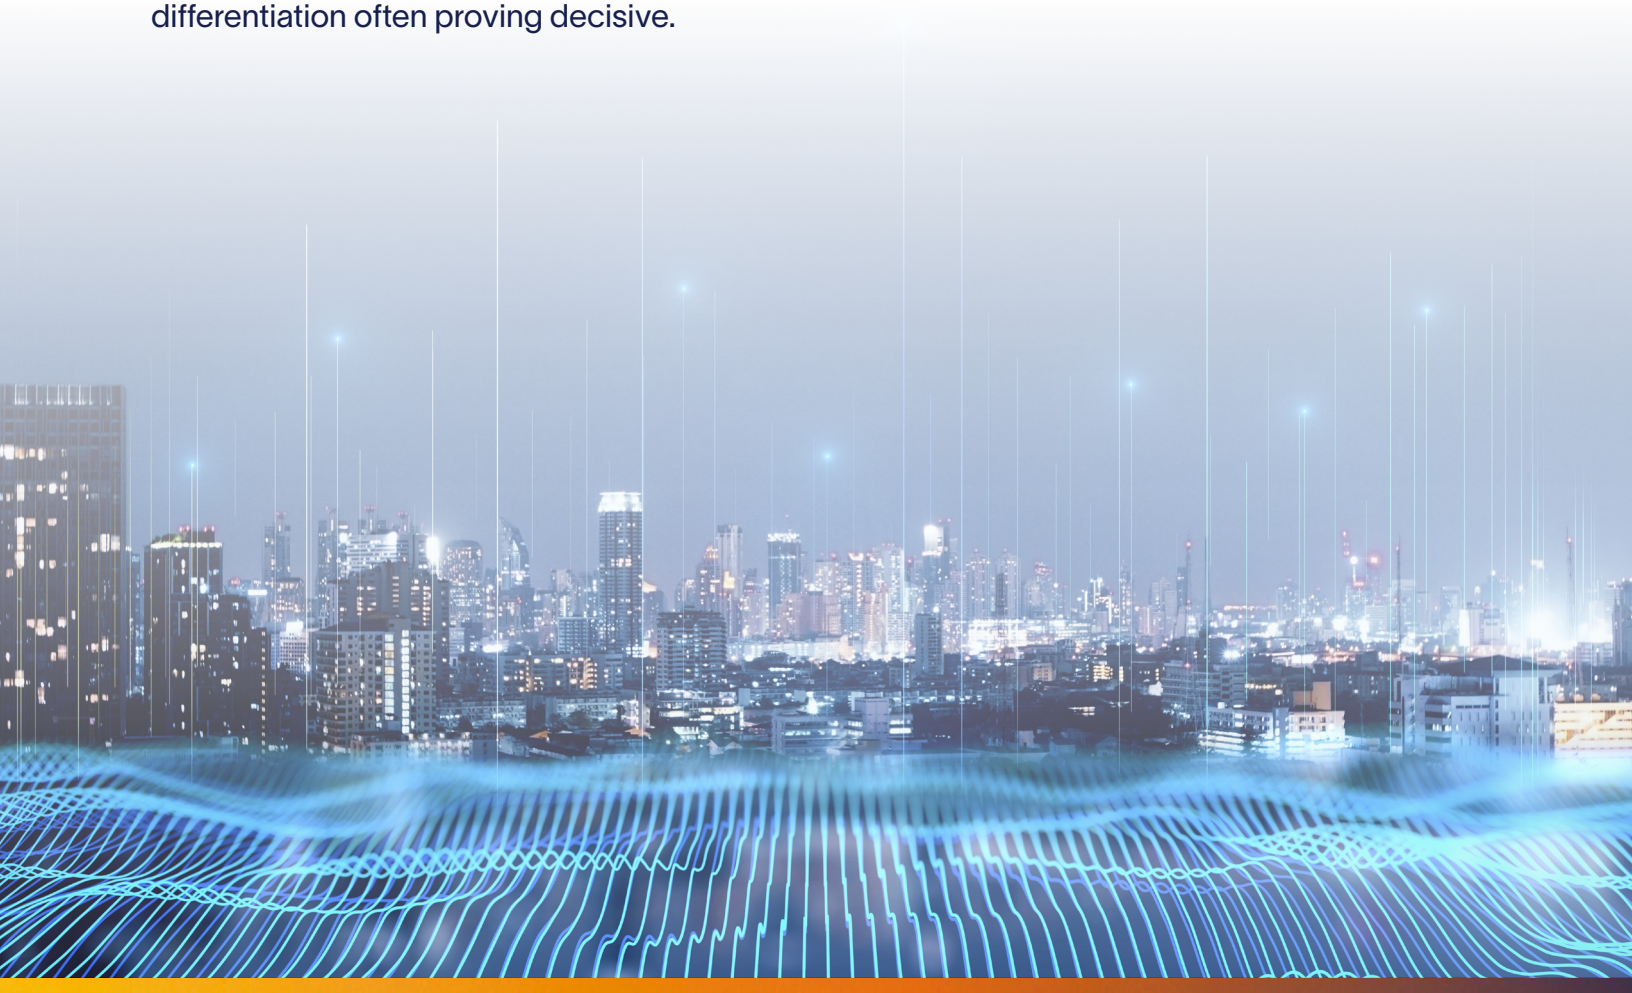

Technical AI resilience

AI reinforces/erodes a provider's security efficacy and technical differentiation

Specialist provider Platform-driven provider

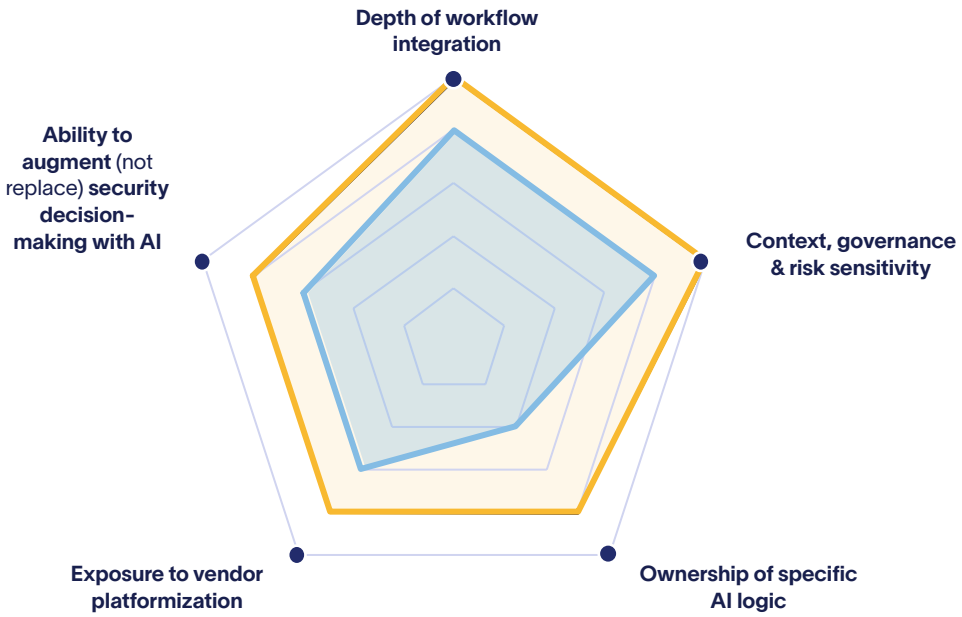

AI tends to reinforce providers whose differentiation is rooted in **context, integration depth, and governance**, while compressing differentiation **based on tooling, templates, or generic automation**

Commercial AI resilience

AI protects/erodes pricing power, trust, and economic control

Specialist provider Platform-driven provider

Cyber Perspective: AI tends to compress differentiation in standalone, tooling-led security products, while reinforcing vendors with deep workflow integration, accountability for outcomes, and trust with buyers

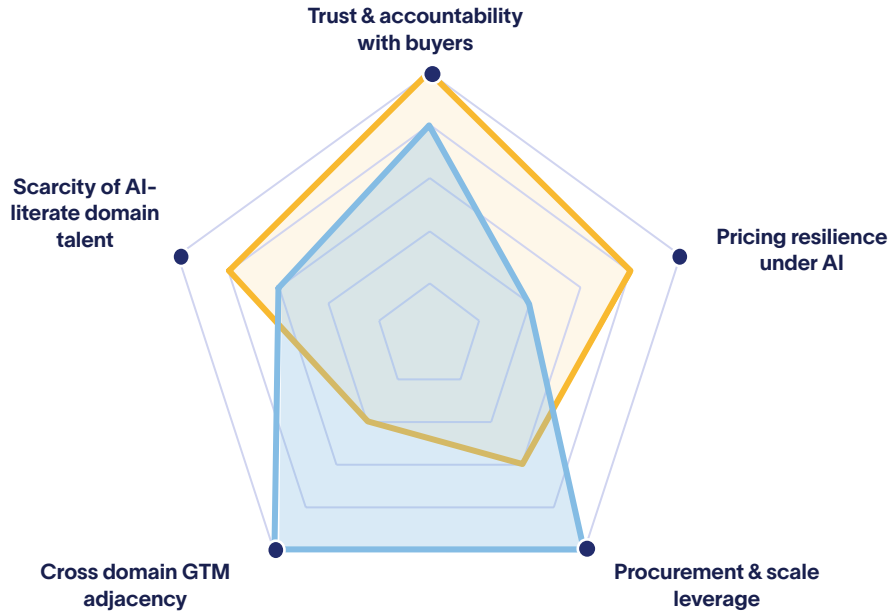

Commercially, **AI often strengthens trust, accountability, and scarcity-based models** while placing pressure on providers whose differentiation **relies primarily on scale or procurement leverage**

Source: Altman Solon

At the company level, technical moats in cybersecurity tend to appear in three main areas:

- **The data**: One of the strongest technical moats comes from large-scale security telemetry collected across many environments. Vendors that aggregate endpoint, network, or identity signals across thousands of customers can build detection and response models that are difficult for new entrants to replicate.
- **The interface & codebase**: Historically a differentiator but increasingly easier to replicate as AI copilots and natural-language interfaces reduce development barriers.
- **Workflow logic**: Differentiation is moving away from rule creation (which AI does well) toward orchestration across multi-vendor environments.

Non-technical moats can matter even more for resilience. In cybersecurity, the "human element" remains a critical barrier to entry. Areas like customer trust, domain knowledge, go-to-market reach, AI-literate talent, and access to capital can reinforce resilience when pure feature differentiation is under pressure.

How can cybersecurity investors apply diligence in an AI-first market?

AI resilience is best understood through an integrated view that combines category exposure, company-level differentiation, and technical validation of product capabilities.

In a market reshaped by AI, competitors can replicate a feature set in months (or weeks) rather than years, raising the bar for what counts as a durable moat. In cybersecurity, differentiation (and outsized returns) increasingly sits in categories anchored in trust, regulated workflows, and proprietary data, and in companies that aggregate signals from multiple systems, guide response actions, and enforce governance.

AI can automate individual tasks like code scanning or alert triage, but replicating these capabilities across complex enterprise environments, with accountable and auditable decision-making, is much harder.

How Altman Solon can help

Altman Solon applies the AI Resilience Diagnostic across the investment lifecycle, both to inform diligence and to translate findings into clear product and growth-value-creation priorities. Typical applications include integrated commercial and technical diligence, anchored in category and company drivers of sustainable differentiation; and pre- and post-deal AI resilience diagnostics that translate findings into product and growth value-creation priorities.

Leadership & Oversight

Ben Matthews
Partner
Ben.Matthews@altmansolon.com

Alexander Jinivizian
Associate Partner
Alexander.Jinivizian@altmansolon.com

Grace Ogilby
Associate Partner
Grace.Ogilby@altmansolon.com

Suhaib Rangoonwala
Associate Partner
Suhaib.Rangoonwala@altmansolon.com

Office Locations

Altman Solon is the leading global strategy firm focused on telecommunications, media, and technology.

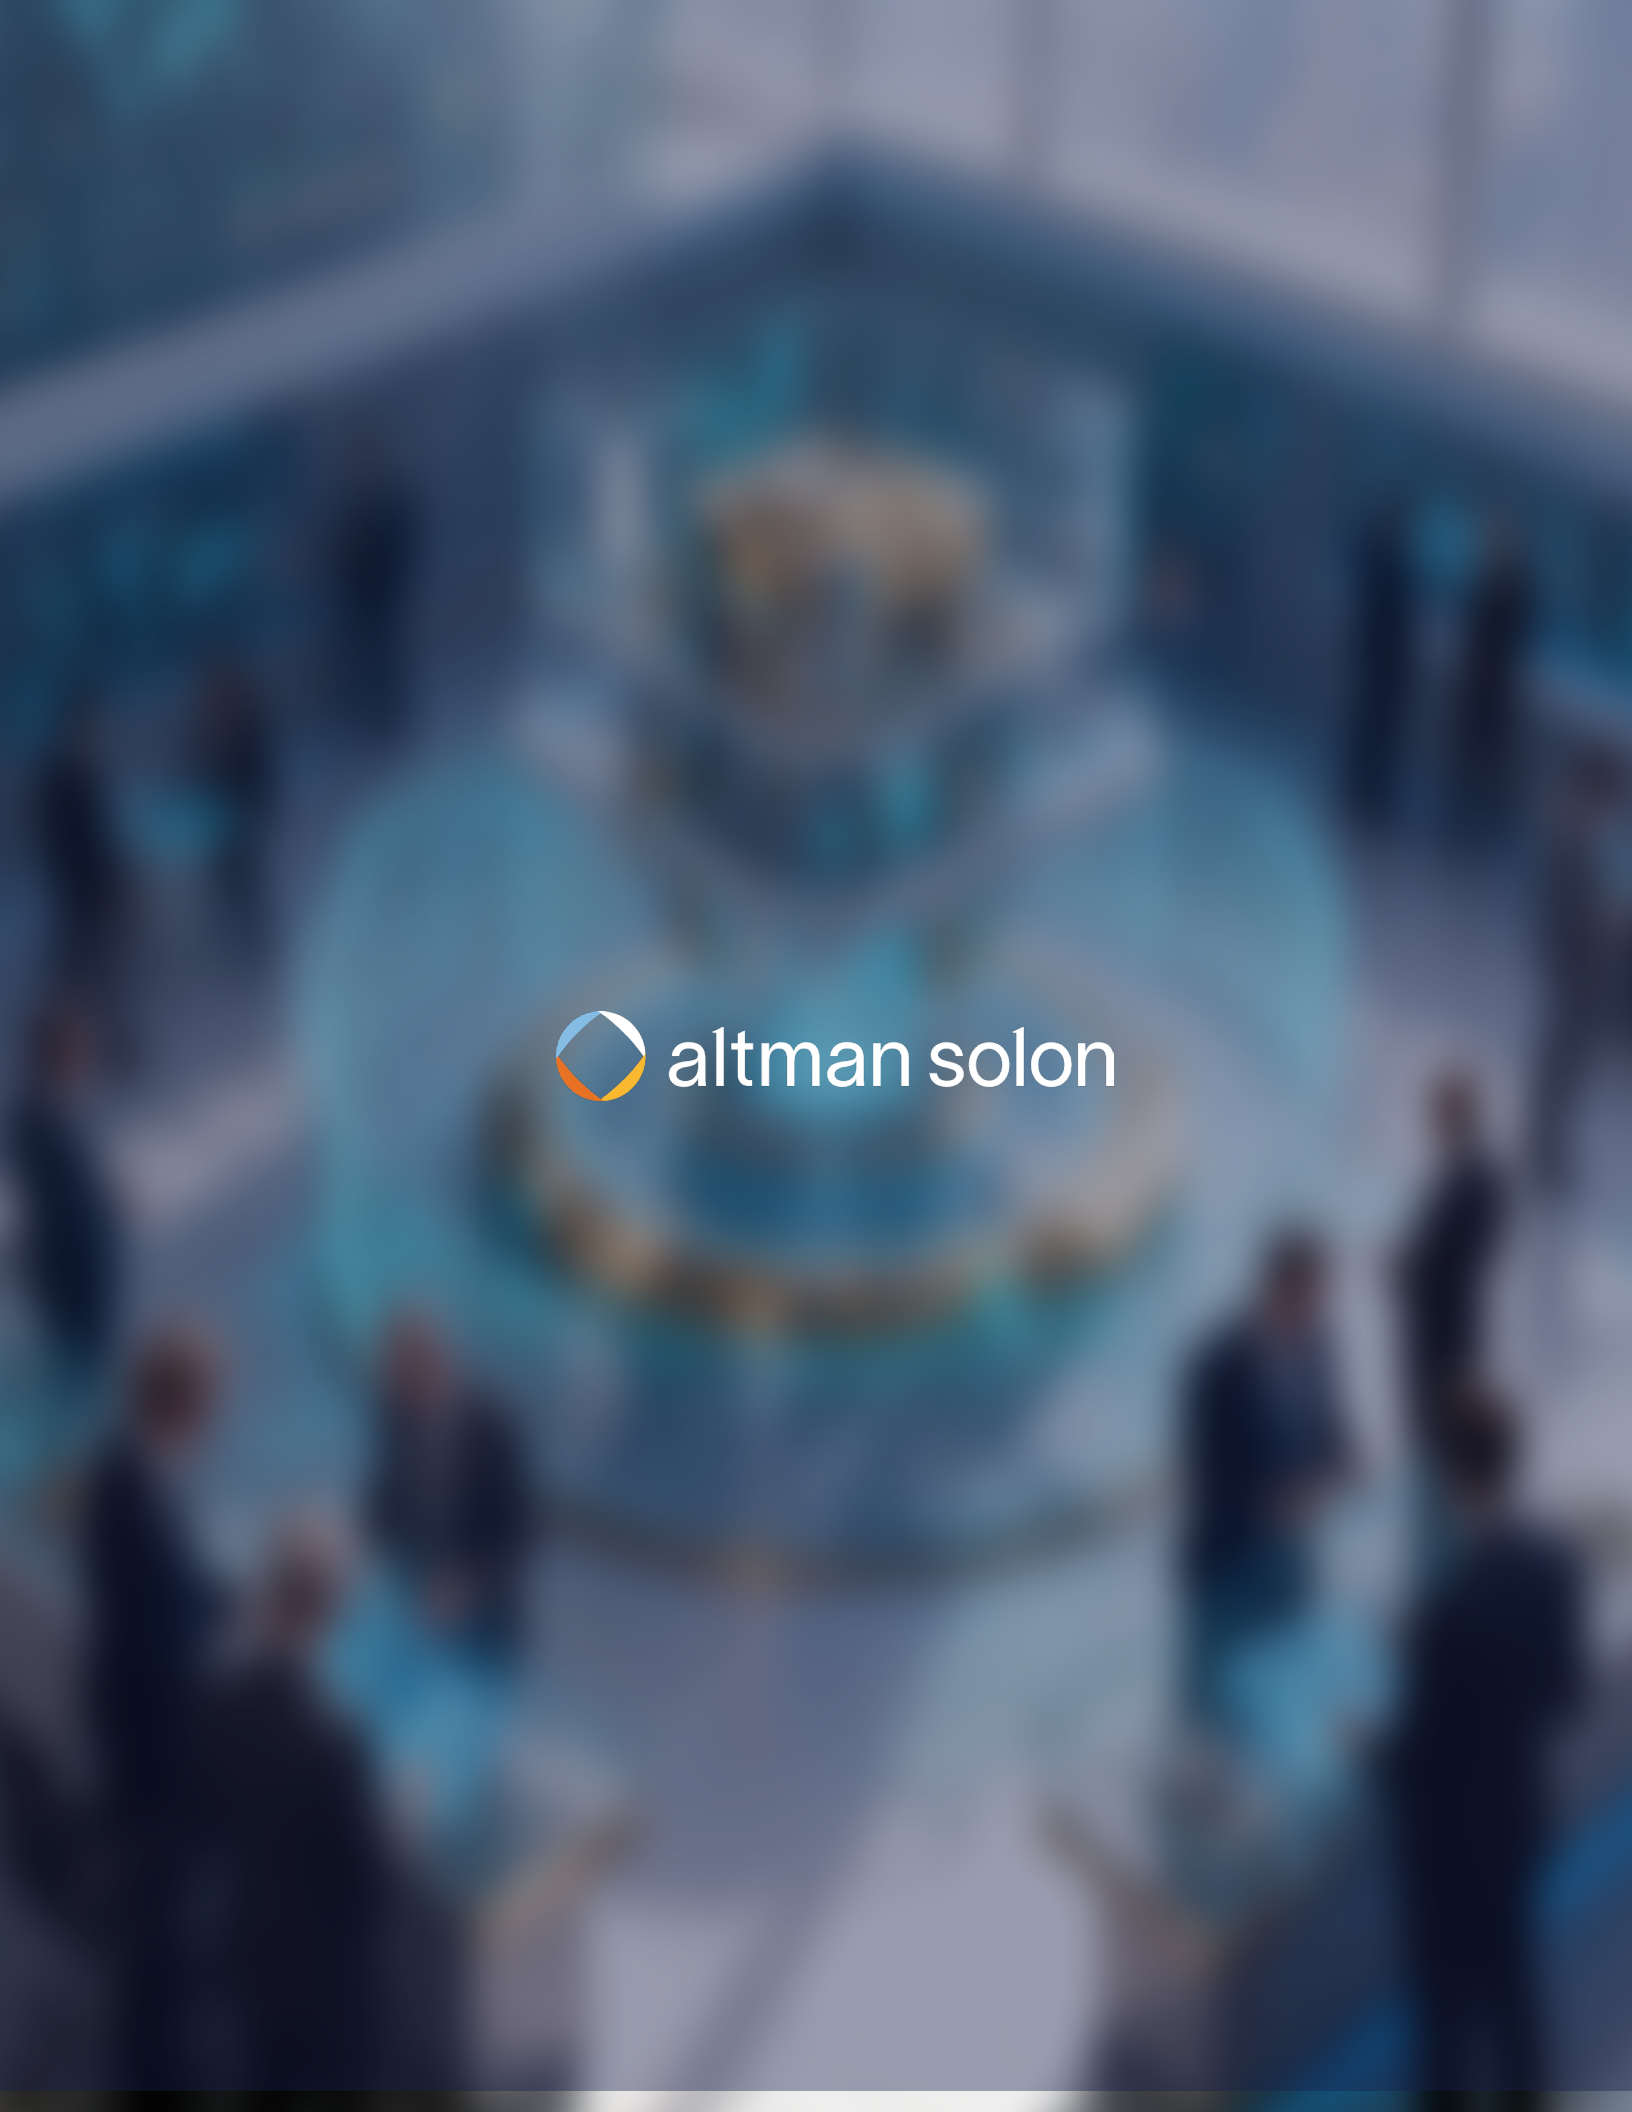
altman solon