



PULSE REPORT

Agentic AI Meets the Enterprise Control Problem



at HUMAN[X] Amsterdam

TABLE OF CONTENTS

Executive Overview 3

CHAPTER 1

From the Lab to the Wild: Reliability Becomes an Infrastructure Problem 4

CHAPTER 2

The Agent Becomes a New Security Principal 8

CHAPTER 3

Human Control Becomes a Risk Decision 13

CHAPTER 4

Context Is the New Control Plane.....18

CHAPTER 5

Escaping POC Hell: AI Value Requires Business Redesign.....23

CHAPTER 6

Trust Is the Scaling Mechanism.....28

Overall Conclusions 33

EXECUTIVE OVERVIEW

The conversation at HumanX Amsterdam was ostensibly about artificial intelligence: models, agents, infrastructure, automation and business value. But across the interviews, another subject kept surfacing underneath it all: control.

As AI moves from answering questions to taking actions, the central enterprise challenge changes from the model's capability to the surrounding organization's ability to constrain, observe, explain and recover from what that model does.

Agentic AI is turning long-standing security principles (least privilege, observability, identity, segmentation, resilience and auditability) into design requirements for an entirely new class of digital actor. The technology is probabilistic, yet it is increasingly being connected to deterministic business systems. It consumes sensitive organizational context.


It receives credentials and tools. It may initiate transactions, alter software, retrieve proprietary data or act over hours rather than seconds.

That makes AI security as much an infrastructure and governance problem as a model security problem.

The strongest consensus across the interviews favored continued AI deployment, and the executives described considerable operational and commercial opportunities. But the organizations likely to capture that value will deliberately engineer boundaries around autonomy, preserve context and provenance, redesign rather than merely automate business processes, and make trust measurable through evidence.

Key Themes

THEME	EMERGING VIEW
Reliability at scale	Error rates that appear tolerable in experiments become operationally significant when agents execute millions of actions. Reliability must therefore be engineered around the model, not assumed from it.
Agent identity and permissions	Agents increasingly resemble digital workers with credentials, tools and authority. Static access control is not enough when a permitted action can still be harmful in the wrong context.
Human control	Human in the loop is becoming risk-dependent rather than universal. Irreversible, subjective and high-consequence decisions attract human judgment; other tasks can operate within bounded autonomy.
Context and provenance	Enterprise context is simultaneously AI's source of value and a major security asset. Organizations need controls over where context lives, how it is retrieved, and whether actions can be reconstructed.
Business redesign	Experiments generate little durable value unless AI is tied to measurable business outcomes and processes are redesigned around new capabilities.
Trust and auditability	Transparency, source verification, privacy and evidence of how decisions were made are becoming prerequisites for enterprise adoption.



CHAPTER 1

From the Lab to the Wild: Reliability Becomes an Infrastructure Problem

The interviews opened with a warning that small-scale AI performance tells enterprises very little about large-scale operational reliability. What looked like a model quality issue in a laboratory becomes an infrastructure, resilience and incident management problem in production.

Robert “Bobby” Blumofe, executive vice president and CTO at Akamai, described scale and reliability as inseparable. At experimental scale, rare failures may never become visible. At enterprise scale, the same underlying rate can generate clusters of failures faster than teams can investigate them.

“Once you get out into the wild in production at scale, now you may be doing millions or billions of interactions, and now you’re going to see many of these failures,” Blumofe said.

If failures multiply with volume, the model is only one component of AI reliability. Jessica Reeves, CEO at Axoniq, argued that dependable AI requires the surrounding workflows, data architecture and distributed systems to be designed for accountability as well. Her concern is particularly acute where legacy architectures preserve current state but discard the historical sequence of decisions that produced it. For autonomous systems, that history becomes evidence.

Reeves prefers an architecture with built-in persistence and provenance, allowing decisions to be reconstructed and replayed. In her account, event sourcing lets you rewind an incident, determine what happened, correlate the relevant elements, and potentially unwind and replay the process.

Logging an agent’s final action will often be insufficient. Investigators may need the preceding tool calls, retrieved context, intermediate state, system conditions and business logic that influenced the decision.



Ivan Burazin
Co-Founder and CEO, Daytona

Infrastructure Built for Humans Meets Infrastructure Used by Agents

Ivan Burazin, co-founder and CEO at Daytona, pushed the scaling argument further. He argues that today's cloud architecture largely assumes human-driven applications and relatively stable software workloads. Agent workloads are different: they repeatedly create, edit, execute and discard working environments and can do so at machine scale.

"Agents will be the number one consumers. They will consume everything," Burazin said.

Whether or not agent populations eventually reach the extraordinary numbers discussed in the interview, the architectural point is relevant now: infrastructure designed around human request rates can behave very differently when software agents become persistent consumers of compute, applications and APIs.

Burazin also distinguished model inference from the work agents perform after reasoning. An agent may use a model to decide what should happen, but it still needs conventional compute to manipulate spreadsheets, run code or execute applications. That expands the infrastructure discussion beyond model-serving GPUs to the general-purpose environments agents use to act.




Jessica Reeves
CEO, Axoniq

Modernization Without a Big-Bang Rewrite

The interviews do not point toward replacing every legacy system before deploying AI. Reeves instead argued for incremental modernization: identify decision paths in existing applications, introduce better decision tracking and provenance, and change architecture without disrupting the entire estate at once.


Her argument is that slowing long enough to build the right accountability foundation can ultimately increase deployment speed, because organizations are not repeatedly rebuilding controls after systems reach production.

The practical distinction for cybersecurity leaders is between AI adoption and AI readiness. An enterprise can connect a model to an API quickly; operating thousands of autonomous actions safely requires a different class of engineering.



CONCLUSION

The interviews frame AI reliability as a systems problem. Model accuracy remains important, but production safety will increasingly depend on provenance, runtime isolation, recoverability, capacity planning and continuous observability. Once scale turns rare errors into routine events, smarter agents are only part of the answer; enterprises also need infrastructure that can contain and explain agent behavior.



CHAPTER 2

The Agent Becomes a New Security Principal

Traditional enterprise access control assumes that granting permission largely answers whether an action is allowed. Agentic systems complicate that assumption.

An agent may legitimately possess permission to read a calendar, query Salesforce, browse the web or modify a development environment. Yet the same permitted action can become unsafe if it is performed for the wrong reason, with the wrong information, or at the wrong point in a workflow.

Cal Al-Dhubaib, principal technologist at Rubrik, argued that AI is magnifying risks already embedded in enterprise environments rather than creating an entirely separate security universe.

“The cost of ignoring the basics has just gotten much higher,” Al-Dhubaib said.

His framing separates the problem into three areas: protecting the enterprise from AI-enabled threats; protecting the AI systems operating inside the enterprise; and using AI to improve defensive operations. Of the three, the rise of internal agents creates a particularly unfamiliar security problem. Al-Dhubaib described them as potential “autonomous insiders”: software systems with tools and legitimate access that can still overreach because of poor context, manipulation or misalignment with the intended task.



Robert Blumofe

Executive Vice President and CTO, Akamai

From Least Privilege to Least Capability

Blumofe connected this new problem to an old cybersecurity principle.

Rather than begin with a highly capable agent and then attempt to prevent unwanted behavior through instructions alone, he advocated designing the agent with only the capabilities required for the task.

“Don’t give it the ability to write code if you don’t need the ability to write code,” Blumofe said.

He described this as least capability, an analog to least privilege. Access controls constrain which systems an identity can reach; capability constraints can also limit the kinds of operations the agent can initiate in the first place.

This is particularly relevant to prompt injection. Blumofe noted that content entering an agent’s context window can change its behavior, including content retrieved through tools from public sources. The more general the agent and the broader its tools, the more paths an attacker - or even unexpected input - can use to influence its actions.



Cal Al-Dhubaib
Principal Technologist, Rubrik

Why Context-Aware Controls Matter

Al-Dhubaib's example of an enterprise search agent illustrates why conventional permissioning is insufficient. An agent may be authorized to access internal Salesforce information and separately authorized to search the public internet. If the agent incorporates sensitive internal material into an external search query, each action may be technically permitted while the combined workflow results in data exfiltration.

His answer is context-based observability and controls: monitoring not only whether an API call is permitted but whether it is consistent with the goal, organizational policy and information context of the task.

Telemetry would have to expand accordingly. Beyond uptime, network traffic and authentication events, security teams may need to monitor the original prompt, retrieved information, tool sequence, model traces and outputs.



Ivan Burazin
Co-Founder and CEO, Daytona

Treat the Agent Like a Digital Worker

Burazin started from a practice enterprises already know well, giving employees role-specific accounts, spending limits and restricted access. In his view, agents can inherit similar patterns rather than being trusted with a user's unrestricted credentials. His personal example involved giving a digital assistant access equivalent to a bookkeeper rather than handing it banking credentials that could transfer funds.


Agents and people are not identical, but identity engineering already offers useful primitives: separate accounts, scoped credentials, approval thresholds, transaction limits and revocation.

The next step is combining those familiar mechanisms with the contextual controls required by probabilistic systems.



CONCLUSION

Agentic AI introduces a security principal that can possess legitimate authority while behaving in unanticipated ways. The interviews converge on a layered response: least privilege, least capability, isolated execution, richer observability and context-sensitive policy enforcement. Security teams should expect identity governance for agents to become as important as identity governance for human users, and potentially more dynamic.



CHAPTER 3

Human Control Becomes a Risk Decision

“Keep a human in the loop” has become a familiar answer to AI risk. The HumanX interviews qualified it.

The speakers generally supported continued human judgment for consequential decisions, but they did not describe human intervention as necessary - or even practical - at every stage of every autonomous system.

The more useful dividing line is reversibility and consequence.

Swami Sivasubramanian, vice president of agentic AI at AWS, described this using the language of one-way and two-way doors. Reversible decisions can be delegated more freely; irreversible actions require stronger judgment and approval mechanisms.

“One-way door decisions are the ones that you cannot reverse back. For that, you need judgment,” Sivasubramanian said.

That framework converts an abstract discussion about autonomy into system design. A low-value, reversible action may proceed automatically. A large refund, major healthcare decision or high-consequence financial action can stop at a human gate.



Patrick Mans

Head, Data and AI Engineering Innovation, Philips

Healthcare Draws a Bright Line

Patrick Mans, head of data and AI engineering innovation at Philips, described a particularly clear boundary in healthcare.

Philips uses AI across imaging and workflow automation, but Mans distinguished predictable machine learning systems used in medical applications from generative systems that can hallucinate. In his description, generative AI helps reduce administrative burden, while clinical decisions retain human authority.

“It’s absolutely the human who gets the final say. Yeah, no doubt about that,” Mans said.

AI-generated reports go to clinicians, who can accept them, change them and feed disagreement back into the system.



Bobby Healy
CEO and Founder, Manna Air Delivery

Physical AI Shows the Opposite Boundary

Bobby Healy, CEO and founder of Manna Air Delivery, illustrated why “human in the loop” cannot be treated as a universal safety mechanism.

“We don’t let the agents fly the aircraft,” Healy said.

Manna uses AI for planning and fleet orchestration, but the flight-control loop itself is deterministic and regulated. Once an aircraft is in flight, however, a remote human may actually be too slow to respond to a fast-moving failure. The safety architecture instead relies on multiple redundant sensors, computers, communication paths and backup systems that can respond automatically.

This distinction is important for cyber-physical systems. Safety may require human authority at the policy and approval layer while automated deterministic mechanisms handle real-time control.



Colin Jarvis

Global Head, Forward Deployed Engineering, OpenAI


Judgment Remains a Human Responsibility

Colin Jarvis, global head of forward deployed engineering at OpenAI, made a similar distinction in enterprise workflows. He described human involvement as particularly important where information is sensitive, judgments are highly subjective, or the consequences of an incorrect action are irreversible.

Soyoung Lee, co-founder and head of go-to-market at TwelveLabs, drew another boundary around interpretation. Video AI can extract scenes, actions and context from large bodies of footage, but she resisted turning that capability into human judgment.


“So, the AI’s job is not to make judgments. It’s just to tell you what it sees,” Lee said.

Security teams using AI to analyze video, alerts or behavioral patterns can apply the same principle, automating perception, retrieval and summarization more aggressively while stopping short of drawing subjective conclusions about intent.



CONCLUSION

Human oversight is evolving from a blanket requirement into an explicit risk control architecture. The interviews suggest that enterprises should classify actions by consequence, reversibility, subjectivity and response time. Humans remain central to high-consequence judgment, while bounded autonomous systems can handle reversible work and deterministic safety mechanisms may outperform human intervention in time-critical scenarios.



CHAPTER 4

Context Is the New Control Plane

AI cannot make strong enterprise decisions from data alone.

Several interviews returned repeatedly to context: market events, internal strategy, conversations, historical decisions, employee knowledge, business processes and the reasons behind earlier actions.

That context is what makes enterprise AI valuable. It also makes enterprise AI dangerous if access is poorly governed.

Bianca Anghelina, founder and CEO at Aily Labs, described context as having multiple dimensions: external market and regulatory factors; forward-looking strategy that cannot be inferred from historical data; and user knowledge that may exist only in people's heads, emails, or conversations.

"So internal data is not enough," Anghelina said.

Her example is instructive. Historical business data might suggest withdrawing investment from an underperforming product, while a pending regulatory event known to the business could reverse the economics within weeks. Data without context can therefore produce a technically rational but commercially wrong recommendation.

Aily's approach embeds business logic in a semantic layer and supplements it with governance, model calibration and ongoing human feedback.



Xuedong D. Huang
CTO, Zoom

Context Becomes a Security Boundary

Xuedong D. Huang, CTO at Zoom, made the security implications explicit.

Frontier models have broad general intelligence, but they don't inherently include a company's internal knowledge, relationships, meetings, decision history or operating logic.

"They don't have the context of enterprise," Huang said.

To make the models useful, enterprises must connect them to exactly that information. Huang therefore described the platform connecting enterprise knowledge and frontier reasoning as a critical security boundary. His preferred architecture includes controlled enterprise indexing and the ability to keep sensitive organizational knowledge within the enterprise while still using external reasoning capabilities.

The risk sits in that connective layer, since the platform that makes an agent useful may also expose an organization's most valuable tacit knowledge.



Umesh Sachdev
CEO and Co-Founder, Uniphore

Do Not Rent the Intelligence That Differentiates the Business

Umesh Sachdev, CEO and co-founder at Uniphore, took that argument into the domain of sovereign AI.

“You shouldn’t rent your intelligence. You should own it,” Sachdev said.

He is concerned that proprietary data and business context increasingly constitute intellectual property. Enterprises therefore need to distinguish ordinary AI workloads from strategically sensitive use cases involving unique processes, underwriting logic, claims handling, marketing methods or other trade-secret-like knowledge.

Sachdev advocated different deployment models by sensitivity: tightly controlled or air-gapped infrastructure for core intellectual property, private environments for less sensitive workloads, and public-cloud services where the risk profile permits.

Organizations that never adopt that precise architecture will still have to decide which intelligence must remain under direct organizational control.




Jessica Reeves
CEO, Axoniq

Provenance Matters as Much as Context

Reeves' emphasis on event sourcing adds a second half to the context problem.


AI systems need historical context to make better decisions, but organizations also need historical context about the AI itself. Capturing decision paths, inputs and business logic creates a record that can support reconstruction, accountability and learning after an incident.

Lee's discussion of video reinforces the point from another direction. Organizations increasingly hold valuable context in recordings and footage rather than structured databases. Her privacy model emphasizes minimizing biometric exposure and, for sensitive customers, moving the model to the protected data rather than moving protected data to the model provider.



CONCLUSION

Context is emerging as both the fuel and the attack surface of enterprise AI. The organizations that capture the greatest value will likely connect models to deeper internal knowledge, but doing so increases the importance of data classification, provenance, architectural isolation and policy controls. Cybersecurity leaders should treat enterprise context as a strategic asset.



CHAPTER 5

Escaping POC Hell: AI Value Requires Business Redesign

Security architecture cannot be separated from business architecture.

If an organization cannot identify what an AI system is supposed to achieve, it cannot rationally decide how much authority to give the system, which data it should see, or how much security investment is justified.

Across the interviews, executives repeatedly rejected experimentation for experimentation's sake.

Sivasubramanian described organizations generating large numbers of demonstrations without defining success criteria, governance or leadership sponsorship. His prescription is to define measurable outcomes, select a small number of projects, and build common infrastructure other teams can reuse, rather than letting each experiment create its own stack.

Tying AI projects to durable business outcomes also makes shared controls easier to justify.



Costi Perricos

Global Generative AI Business Leader, Deloitte

The Cappuccino Problem

Costi Perricos, Deloitte's global generative AI business leader, argued that many organizations are aiming too low.

Improving every step of an existing process by a few minutes can create efficiency without fundamentally changing economics.

"You free up 10 minutes here and 10 minutes there, and your employees can go have a cappuccino, but it never really hits the hard line," Perricos said.

He called this "cappuccino ROI." His alternative is to redesign the process around what AI makes possible and then redistribute work between digital and human participants.

That redesign introduces its own control problem. Perricos noted that humans may transition from creators to auditors and checkers, roles that require different skills and can introduce new error patterns when people are disengaged from the underlying work. Human oversight therefore must be designed, not simply added.



Colin Jarvis

Global Head, Forward Deployed Engineering, OpenAI

Start With the Business Lever

Jarvis reached a similar conclusion from deployment work.

“What are the biggest levers I would pull in my business if I wanted to change my business?” Jarvis asked.

His answer is to start with the business problem rather than start with an AI capability and look for somewhere to insert it.

That principle also affects security architecture. A strategically important workflow can justify strong governance, data controls and approval systems. A trivial use case may not. Risk should follow the value and consequence of the process being redesigned.



Parag Parekh

Chief Digital Officer, IKEA

Measurement and Adoption Are Different Problems

Parag Parekh, chief digital officer at IKEA, described a deliberately measurable AI strategy spanning leadership awareness, usage frameworks, technology-sourcing decisions and prioritized use cases. IKEA evaluates both growth and cost effects and attempts to translate productivity gains into financial impact.

At the same time, global deployment reveals a different bottleneck: people.

“It’s no longer technology for the sake of technology. It’s about how it’s helping them in their day-to-day,” Parekh said.


Rolling a centrally built digital capability into multiple markets can be relatively straightforward. Embedding a new way of working across more than 160,000 employees is harder, Parekh said, making adoption and usage metrics important indicators alongside financial ROI.

Governance that people do not adopt drives workarounds instead of compliance, which is why usable, trusted AI controls are more likely to become operational ones.



CONCLUSION

The interviews reject the idea that AI value emerges simply from making existing tasks faster. Durable value comes from identifying consequential business outcomes, redesigning work around AI's capabilities, measuring the results and changing how people participate. Security and governance belong in the design of that new operating model from the start.



CHAPTER 6

Trust Is the Scaling Mechanism

Trust is sometimes presented as a soft concept: something earned through communication, brand or policy. In the interviews, trust was far more operational.

It meant knowing what an agent is permitted to do. It meant retaining evidence about what happened. It meant being able to inspect sources, protect privacy, understand context, recover from error and decline use cases that violate organizational values.

Stefan Weitz, CEO and co-founder of HumanX, said conversations with large, regulated enterprises revealed far more centralized control over AI than the industry's experimentation narrative might suggest.

"A very tight leash on where it can be used and for what," Weitz said.

He attributed that posture to governability, auditability, transparency and trust. The observation helps explain why impressive demonstrations do not automatically translate into enterprise-scale deployment. The more consequential the organization, the more evidence it needs that AI can operate within policy.



Jim Lanzone
CEO, Yahoo

Trust Also Requires Visible Sources

Jim Lanzone, CEO at Yahoo, approached trust from the information ecosystem.

AI search can synthesize complete answers without requiring users to visit original sources, but that convenience can obscure provenance and make stale or hallucinated information difficult to detect. Lanzone argued for visible links back to publishers, both to help sustain the open web and to let users validate where information originated.

“So, trust but verify is important,” Lanzone said.

Inside the enterprise, AI-generated output should not become authoritative merely because it is fluent, and high-consequence decisions need evidence chains.



Parag Parekh
Chief Digital Officer, IKEA

Responsible AI Includes Saying No

Parekh's account of IKEA provides another dimension of trust: organizational restraint.

IKEA has defined categories in which it is willing to experiment and others in which it is not. Parekh described shutting down experiments when they conflicted with the company's framework, including use cases that could encroach on customer or co-worker privacy and potential applications affected by fairness concerns.

Governance covers more than controls around deployed technology; it includes the authority to decide that a capability should not be deployed at all.



Jessica Reeves
CEO, Axoniq


Auditability Becomes the Receipt

For Reeves, the regulatory and accountability requirement is straightforward: organizations should be able to demonstrate what happened, why it happened, and what caused the result.

“You have to bring your receipts, if you will,” Reeves said.

That phrase captures much of the HumanX discussion.

Organizations can move quickly with AI, but enterprise trust ultimately depends on evidence. Producing it requires systems observable enough to investigate, authority narrow enough to contain failure, controlled context, attributable decisions, and an appropriate human or deterministic gate for high-consequence actions.



CONCLUSION

Trust is the outcome technical control is intended to produce. Enterprise AI will scale where organizations can show - not merely claim - that systems are governed, observable, accountable and aligned with the sensitivity of the work they perform.

OVERALL CONCLUSIONS

The HumanX interviews collectively suggest that the industry's center of gravity is shifting from model capability to operational control.

For cybersecurity leaders, six priorities follow:

1. Treat agent infrastructure as production infrastructure. Design for rare failures occurring at massive scale, not for the behavior seen in small demonstrations.
2. Give agents identities, not borrowed human credentials. Combine least privilege with least capability, scoped accounts and transaction limits.
3. Make autonomy risk tiered. Require human judgment for irreversible, subjective and high-consequence decisions while allowing bounded autonomy for reversible work.
4. Protect context as intellectual property. Enterprise conversations, workflows, history and tacit knowledge can be more sensitive than the model itself.
5. Instrument the full decision chain. Prompts, retrieval, tools, context, outputs and actions increasingly belong in the observability model.
6. Demand measurable business outcomes. AI governance becomes more durable when deployments are tied to explicit value and deliberately redesigned processes.

The common thread is architectural discipline. HumanX's interviewees were describing what must exist around AI before enterprises can safely grant it meaningful authority.

The competitive divide may therefore be less about who has access to the most capable model and more about who can turn probabilistic intelligence into a controlled, observable and accountable operating system for the enterprise.

About ISMG

Information Security Media Group (ISMG) is the world's largest media organization devoted solely to cybersecurity, information technology, artificial intelligence and operational technology. Each of its 38 media properties provides education, research and news that is specifically tailored to key vertical sectors including banking, healthcare and the public sector; geographies from North America to Southeast Asia; and topics such as data breach prevention, cyber risk assessment, OT security, AI and fraud. Its annual global summit series connects senior security professionals with industry thought leaders to find actionable solutions for pressing cybersecurity challenges.

Contact

(800)944-0401 • sales@ismg.io

