VIEWPOINT

SECURING THE AI-FIRST WORKPLACE: RETHINKING TRUST FOR A MODERN WORKFORCE OF PEOPLE, AI AGENTS, AND ROBOTS

Infosys
Navigate your next

Executive Summary

The AI-first workplace is transforming how enterprises operate as people, AI agents, and robots increasingly work together across digital and physical environments. This shift expands the enterprise attack surface and exposes the limits of security models designed primarily for human users and managed devices. As work increasingly flows between people and intelligent systems, security must evolve from protecting individual assets to governing digital interactions across the enterprise. Securing this new operating model requires stronger AI governance and cyber resilience built into everyday operations, continuous validation of trust, and privacy embedded throughout the AI lifecycle.

The Workplace Has Entered the AI-First Era

Enterprise work is evolving. AI is moving beyond isolated employee assistance as organizations deploy intelligent systems across customer service, software development, finance, manufacturing, and operations. At the same time, productivity assistants are becoming autonomous agents that can make decisions and act with limited human oversight. As AI becomes embedded in business processes, employees increasingly collaborate with intelligent systems in their daily work, fundamentally reshaping enterprise workflows.

The Siemens–NVIDIA partnership illustrates how AI is becoming part of enterprise operations. By combining AI-powered digital twins — virtual representations of physical systems — with industrial AI and robotics, manufacturers can simulate production before deployment and make more informed operational decisions. Although rooted in manufacturing, the collaboration reflects a broader shift toward people, AI systems, and intelligent machines working together across shared enterprise processes.[1]

Every AI agent, robotic system, cloud service, partner connection, and machine identity expands what enterprises must secure. As work flows across people, AI agents, and machines, traditional security assumptions become harder to maintain.

Security for the AI-First Workforce

Enterprise security has evolved considerably over the past two decades. Identity and Access Management (IAM), cloud security, and endpoint security have significantly strengthened enterprise defenses, while Zero Trust Architecture has shifted security toward continuously verifying users and devices instead of trusting them by default. Enterprises had designed most of these capabilities for environments where people were the primary users. Now, AI agents increasingly retrieve information and initiate enterprise workflows, while machine identities and robots interact directly with enterprise systems. Authentication remains essential, but it is no longer a fool-proof basis for trust.

As a result, cyber risk has become more interconnected. As non-human identities participate in enterprise operations, a compromised AI agent or machine identity can influence multiple applications and workflows. Security now supports business continuity as much as it protects enterprise systems.[2] Organizations therefore need a security model that continuously governs identities across people, AI agents, and robots, something a patchwork of separate controls cannot achieve.

Understanding the Cyber Risk Trifecta

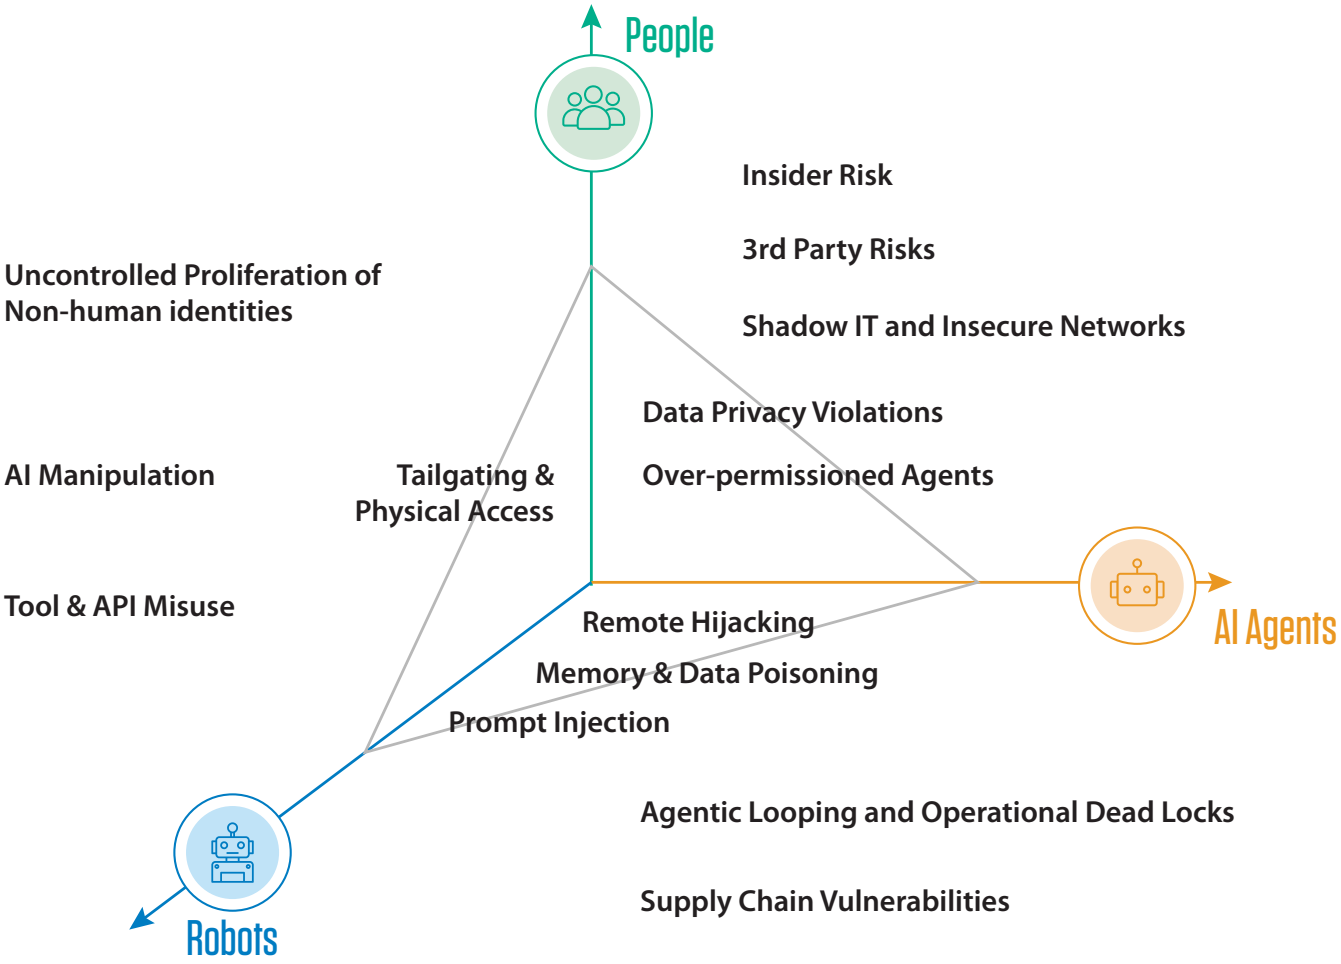

We have observed a Cyber Risk Trifecta evolving with enterprises which enables us to focus on the relationships between people, AI agents, and, in some environments, robots. Every handoff introduces a decision across four areas[3]:

- Identity (Both Human and Non-Human Identities)
- Authority
- Permissions
- Accountability

Traditional security models evaluate users, devices, and applications as individual entities. In the AI-first workplace, however, security decisions span multiple participants. An employee may assign work to an AI agent, which then directs a robotic system to complete the final task. Even when no participant is compromised, failures in permissions, context, or oversight can disrupt the workflow and blur accountability.

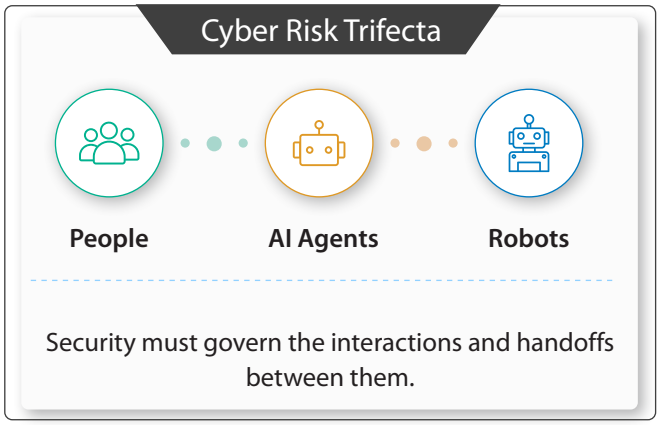

These failures make it harder to determine who is responsible when work is shared across employees, AI agents, and automated systems. Security has to govern the entire workflow rather than each participant in isolation.

From Framework to Action

Understanding the Cyber Risk Trifecta is only the starting point. Organizations also need a practical approach to securing the AI-first workplace without slowing AI adoption.

The framework rests on three priorities:

Cyber Resilience
Keep critical operations running through disruption.

Govern Identity
Govern AI agents and machine identities throughout their lifecycle.

Continuous Trust Model
Evaluate interactions and actions continuously as conditions change.

Cyber Resilience

For organizations operating in an AI-first workplace, resilience has become a business requirement rather than a security objective. In such interconnected environments, preventing every cyber incident is no longer a realistic goal. What matters is limiting disruption and keeping critical operations running.

The FIFA World Cup and other such large international events offer a useful example.[4] The technology supporting millions of viewers operates under constant cyber pressure, yet the expectation is uninterrupted service. Success depends on preparation, rapid response, and the ability to recover without affecting the participants' or spectators' experience. Organizations across industries face the same challenge: critical operations cannot stop because a cyber incident occurs. That is why resilience should be built into day-to-day operations rather than treated as a specialist security function.

- Test critical workflows
- Rehearse recovery scenarios
- Maintain AI-enabled operations during disruption

Govern Non-Human and Human Identity

Identity has always been at the center of enterprise security. What has changed is the range of identities that now participate in the workflow. Enterprise identities now include AI agents, machine identities, APIs, service accounts, applications, and robotic systems, many of which operate autonomously and at machine speed.[5]

Infosys' work on securing the growing digital workforce highlights that securing non-human identities requires more than just authentication. Clear ownership, continuous monitoring, governance, and Zero Trust are essential as AI agents and machine identities take on larger roles in enterprise operations.[6]

Managing these identities isn't just about just confirming who or what is requesting access. Non-human identities deserve the same level of oversight as human users because they are increasingly performing work with a direct business impact.

- Define roles and ownership
- Govern permissions throughout the lifecycle
- Monitor behavior for unusual activity

Continuous Trust Model

After a user logs in at an AI-first workplace, trust has to be established continuously rather than assumed. Every interaction involving people, AI agents, robots, and connected systems should be evaluated in context.

Zero Trust established the principle of continuous verification for users and devices. The AI-first workplace extends that principle to autonomous software, machine identities, and robotic systems. Access decisions increasingly depend on context as much as credentials, requiring organizations to determine not only who initiated an action but whether it is appropriate under the current business and operational conditions. Those same contextual signals must also be governed carefully to protect personal information and meet privacy obligations.

Continuous visibility relies on identity verification and behavioral analytics to detect unusual activity. These insights become more effective when combined with AI security posture management, threat intelligence, threat exposure management, and AI-driven incident response. Organizations should review governance policies regularly and ensure resilience exercises reflect realistic operating conditions. Security controls should also adapt as AI capabilities mature. Security becomes an ongoing discipline through AI-driven security operations that support changing business processes.

- Verify identity and behavior continuously
- Evaluate actions in context
- Adapt controls as AI capabilities evolve

What Comes Next for Workplace Security

Without these capabilities in place, security gaps will only widen as AI agents and robots take on more enterprise responsibility, leaving trust decisions fragmented across disconnected systems and functions. The next phase of workplace security requires these capabilities to evolve alongside the workforce itself, with resilience, identity governance, and continuous trust built into everyday operations rather than treated as a specialist function.

For Infosys, securing the growing digital workforce depends on bringing together cyber resilience, identity governance and continuous trust, across the digital workplace. Organizations that embed these capabilities into everyday operations will be better positioned to build a secure digital workplace that supports innovation without compromising security, enterprise trust or user confidence.[7]

References

1. https://press.siemens.com/global/en/pressrelease/siemens-and-nvidia-expand-partnership-build-industrial-ai-operating-system
2. https://www.infosys.com/services/cyber-security/documents/growing-digital-workforce.pdf
3. Infosys CyberNext Workplace Security and Privacy
4. https://www.weforum.org/stories/cybersecurity/cybersecurity-digital-risk-fifa-world-cup/
5. https://www.infosys.com/iki/technology-review/deep-tech-perspectives/behavioral-fingerprinting-nonhuman-identity-security.html
6. https://www.infosys.com/services/cyber-security/documents/growing-digital-workforce.pdf
7. https://www.infosys.com/services/cyber-security/documents/growing-digital-workforce.pdf

About the Author

Karthik Andhiyur Nagarajan

Practice Manager and Senior Industry Principal, Infosys

Karthik heads Cyber Platforms and AI & Data Protection Practice at Infosys. He has 20+ years of experience in Product design and consulting. He is an expert on AI, Data Protection and Customer Experience Strategy.

For more information, contact askus@infosys.com

Infosys
Navigate your next