Microsoft

AI Agents and Enterprise Identity Management: Filling the Governance Gap

Tech Mahindra & Microsoft Perspective on Securing Autonomous Systems through Adaptive Trust

FEATURING RESEARCH FROM FORRESTER

Closing The AI Trust Gap With Service Providers

EXECUTIVE SUMMARY

Enterprise AI agents have matured from prototyping to production. They are now making decisions that are critical to the enterprise, like approving transactions, allocating resources, and accessing sensitive data, often with minimal human supervision. This shift exposes a fundamental gap: existing identity and access management frameworks, built for predictable, rule-based operations, cannot govern nondeterministic AI systems that think, act, and adapt autonomously.

The scale is already visible in reports. For instance, 80% of organizations report that their AI agents have taken unintended actions, including accessing unauthorized systems, sharing sensitive data, and downloading restricted content, according to SailPoint, an enterprise identity security company.[1] Despite this risk, Delinea, an identity security and privileged access management provider, observed that only 55% of organizations have access controls in place for AI agents.[2]

Against this backdrop, this study analyzes why traditional identity access management (IAM) fails with AI agents and how enterprises can close that gap, drawing on OWASP's GenAI Security recommendations as the governing framework. Microsoft brings forward the identity and security infrastructure through Entra Agent ID, Purview, Defender, and Foundry. We complement this with our orchestration and validation capabilities through our agentic AI platform, Orion, and our AI validation and verification solution, VerifAI. Together, the joint stack addresses the governance gap from both ends: platform-level identity controls and implementation-layer behavioral validation.

IN THIS DOCUMENT

FROM BOTS TO BRAIN: THE AUTONOMY SHIFT

When an AI agent makes a wrong call like approving a fraudulent transaction, misallocating resources, or exposing sensitive data, who is accountable? With robotic process automation (RPA), the answer was straightforward: a line of code, a process owner, or a system administrator.

With AI agents, the tables have turned. These systems make autonomous decisions, operate at speed, and do not follow a deterministic path. Consider the difference: a procurement bot fills out a form, with every step known and codified. A procurement agent forecasts demand, maps cost against delivery speeds, factors in vendor quality history, and then commits capital. Same function, completely different risk profile. This probabilistic way of operating makes it difficult to bind agents in a traditional responsible, accountable, consulted, and informed (RACI) matrix.

Traditional IAM is built for static roles. AI agents are anything but. They follow or create execution paths that no human programmed.

REAL WORLD CHALLENGES

As agentic AI systems move into enterprise environments, their autonomous decision-making and multi-step execution introduce challenges that traditional systems cannot manage. They operate across tools, data, and workflows with limited predictability, creating gaps in accountability, access control, and governance:

- **Ambiguity in Accountability:** When an RPA bot falters, responsibility typically falls on the process owner. With AI agents, accountability is unclear. For example, an insurance agent may incorrectly deny a claim due to a combination of ambiguous data, model behavior, and intermediate decisions. Determining responsibility requires getting into the 'why,' not just 'what,' across the decision chain.

- **Accelerated Permission Creep:** Permission creep (gradual accumulation of access rights) builds slowly in legacy systems, with each access change recorded and approved by humans. In agentic AI-driven environments, however, permissions can expand rapidly. There are emerging cases of agents capable of autonomously gaining access to multiple systems, increasing the risk surface.

- **Regulatory and Legal Uncertainty:** Regulatory frameworks have not yet caught up. Organizations are still navigating uncharted territory where AI agents make logically sound but functionally wrong choices, leading to financial and reputational losses.

There is no silver bullet. In this case, overcontrol will kill innovation and value, while insufficient controls will cause damage and liability.

WHY OLD IAM DOESN'T MEET THE RISKS OF AGENTIC AI

While the challenges above expose operational gaps, traditional IAM's core foundations do not hold for agentic systems. Designed for humans and scripted automation, IAM presumes stable, observable actors.

Agents upend this with their reasoning, adaptation, and emergent behaviors:

- **Fixed Identities:** Traditional IAM expects access needed to remain steady. Agents discover new integrations on their own; if connecting to another system improves performance, they may autonomously establish those connections. Their requirements shift constantly.

- **Routine Patterns:** Security models bank on familiar behaviors. Agents do not comply. They forge paths that might be innovative but unfamiliar to the humans governing them.

- **Clear Ownership:** Governance assumes clear ownership. With agents serving multiple teams, making decisions that ripple across different risk profiles and regulatory domains, assumptions fall apart fast.

THE EMERGING THREAT LANDSCAPE

The gaps in IAM described above open active attack vectors. It is no longer about prompt injections. Enterprises now face goal hijacking (manipulating what an agent is trying to achieve), memory poisoning (corrupting the context an agent carries across sessions), and cascading failures across agent clusters. These attacks exploit the very capabilities that make agents valuable: their autonomy, memory, tool access, and coordination.

To bring structure to this evolving risk landscape, the OWASP GenAI Security project released the Top 10 for Agentic Applications 2026 in December 2025, establishing the industry benchmark for agentic AI security. [3]

The table below maps select OWASP risks as a snapshot of our joint response with Microsoft. While these risks continue to evolve, mitigation capabilities are being developed across both platforms.

OWASP Risk	Description	TechM-Microsoft Mitigation
ASI01: Agent Goal Hijack	Malicious content manipulates agent objectives	VerifAI validation (TechM), Purview prompt filtering (Microsoft)
ASI02: Tool Misuse	Agents use legitimate tools destructively	Orion tool (TechM) permission scoping, Foundry (Microsoft) behavioral policies
ASI03: Identity & Privilege Abuse	Agents operate with excessive or compromised permissions	Entra Agent ID, Conditional Access, Just in time (JIT) tokens (Microsoft)
ASI04: Unexpected RCE	Agents execute unvalidated code	Foundry sandboxed execution, Defender threat detection (Microsoft)
ASI05: Supply Chain	Agents rely on malicious dependencies or compromised tools	Purview classification, Foundry tool restrictions (Microsoft)
ASI06: Memory Poisoning	Corrupted context persists across sessions	VerifAI memory lineage tracking (TechM), Session isolation (Microsoft)

Figure 1: OWASP Agentic AI Risk Landscape and Joint Mitigation Approach

Our TechM-Microsoft stack addresses these risks end-to-end: Entra for identity, Purview for governance, Defender for threat detection, and Orion and VerifAI layers for orchestration security and behavioral validation, respectively.

BUILDING A FLEXIBLE TRUST SYSTEM

Enterprises need a framework built for how agents actually operate. The following principles for adaptive trust are distilled from our proven real implementation experience. Tech Mahindra Agentic AI implementations adhere to the critical requirements in enterprises along with support from Microsoft offerings including:

- **The 'Least Agency' Principle:** Agents should get only the minimum autonomy needed for bounded tasks: restricting actions (not just data), setting thresholds that trigger human approval, time-limiting elevated permissions, and capping the number of sub-agents that can be spawned
- **Intent-Based Authorization:** The agent's intentions need to be authenticated. Entra Agent ID does this through scoped JIT tokens (specific file, specific channel, not blanket access), Conditional Access evaluated against risk score and behavioral baseline, and action validation before execution rights are granted.
- **Spotting Behavioral Shifts:** Each agent builds its own 'fingerprint' that covers its decision speed, data habits, and thinking routes. VerifAI tracks these baselines and flags deviations.

HOW THE TECH STACKS UP: TECH MAHINDRA & MICROSOFT

These governance principles are only as strong as the architecture behind them. Let us look at how our partnership has built those controls into a stack that enterprises can begin deploying today.

By leveraging the recent capability release by Microsoft including Agent 365, Entra AgentID Tech Mahindra aims to secure the enterprise agents we build for our customers against the various challenges already highlighted above.

With these capabilities of Microsoft, Tech Mahindra can assist its customers to address governance gaps that have challenged enterprise AI deployments:

- **Shadow AI and Agent Sprawl:** The Agent Registry and Entra Agent ID ensure every agent created, whether sanctioned or unsanctioned, is registered, governed, and audited
- **Lack of Agent Identity:** Entra Agent ID provides each agent with a persistent, role-based, auditable identity
- **Inconsistent Access Control:** A least-privilege policy is enforced by default, with Conditional Access applied dynamically to agent activities across the enterprise
- **Limited Observability into Agent Decisions:** Foundry's observability capabilities enable continuous monitoring of agent traces, decisions, and behaviors in real time.

IMPLEMENTATION FRAMEWORK

Drawing from our implementation experience, we recommend the following framework (table below) for agentic AI security, aligned with OWASP guidelines:

Pillar	Objective	TechM-Microsoft Implementation
Discover and Classify	Inventory all agents across hybrid/ multi-cloud; classify by sensitivity, privileges, and business impact	Entra Agent Registry for unified discovery (Microsoft), Orion for cross-cloud visibility, classification based on data access, and action capabilities (TechM)
Define Roles and Guardrails	Set clear operational boundaries for each agent classification; policy-based access tied to specific tasks	Agent Blueprints in Entra (Microsoft), Foundry behavioral policies (Microsoft), VerifAI rule definitions aligned to business intent (TechM)
Enforce Least Privilege and Just-in-Time	Replace standing privileges with JIT access; grant only what is needed, when needed	Entra Conditional Access for agents (Microsoft); scoped token requests (Microsoft); Orion dynamic permission orchestration (TechM)
Authenticate by Intent	Require verifiable identities for all agent interactions; validate that actions match approved use cases	Entra Agent ID credential-free authentication (Microsoft); VerifAI action validation (TechM), and Purview data access governance (Microsoft)
Monitor, Detect, and Improve	Continuous behavioral monitoring; anomaly detection; cryptographic logging; regular security testing	Foundry telemetry and observability (Microsoft); Defender threat detection (Microsoft); VerifAI drift detection (TechM); AI Red Teaming

Figure 4: Five-Pillar Agentic AI Security Framework

GUIDANCE FOR LEADERS

AI agents are already in production, and the governance conversation cannot wait. The priorities, the risks, and the decisions look different across leadership functions:

- **Tech Chiefs (CIOs/CTOs):** Rethink governance from the ground up. Traditional IAM solutions will break when managing AI agents that reason autonomously. As a result, gaps surface quickly at scale. Start building adaptive trust now; bolting it on later is messy and expensive.
- **Risk and Compliance Pros:** Update existing playbooks to account for agentic AI-specific risks for 'reasoning drift' and chain-reaction failures across agent clusters. These are showing up in live deployments. Adopt tools like VerifAI to validate agent behavior at AI speed while keeping audits airtight.
- **Business Executives:** Carefully calibrate governance controls to the organization's risk appetite. Rigid controls kill the speed advantage agents are deployed to deliver. Too little oversight, and the organization is answerable to the board for incidents that could have been avoided. Strike the right balance early; it is a hard step toward succeeding with AI agents.

We foresee that over the next year, AI governance tools will merge into all-in-one suites. Regulations brewing in the EU, US, and APAC will mandate controls for autonomous tech. Business leaders should jump in early and adapt, as waiting for mandates and expiry dates will lead to a scramble.

WRAPPING UP

The shift from experimental to enterprise-grade AI has happened faster than most governance frameworks anticipated. Agents are making decisions in production today without proper reins like identity controls, access boundaries, or audit trails. The real question: how to govern them in a way that unlocks, not undermines, their potential?

Legacy IAM fails with agents because it expects the predictable, while AI agents are built to manage the unexpected. That gap is where governance breaks down and closing it requires a fundamentally different approach.

We bring implementation depth through Orion, our agentic AI platform, while Microsoft brings the identity, security, and compliance infrastructure. Together, the joint stack addresses what neither side could possibly solve alone.

Enterprises entering the Agentic AI journey have a clear choice: build strong governance now with trusted allies or discover the hard way why legacy IAM cannot govern autonomous systems.

The right infrastructure, implementation experience, and governance posture are available today. The only variable is how quickly organizations choose to act.

ABOUT OUR ALLIANCE

Tech Mahindra has 30+ years in enterprise tech. Orion and VerifAI are how we operationalize agent governance.

Microsoft provides the platform—Entra, Purview, Defender, and Foundry. Our engineering collaboration and customer implementations are a testament to the complex issues resolved with the joint approach.

END NOTES

1. (2025, February 12). SailPoint AI agent adoption report. SailPoint Technologies Holdings, Inc. https://www.sailpoint.com/press-releases/sailpoint-ai-agent-adoption-report

2. (2025, September 03). Delinea report reveals only 44% of organizations are fully equipped to support secure AI. Illinois Tech Journal. https://www.illinoistechjournal.com/article/845658152-delinea-report-reveals-only-44-of-organizations-are-fully-equipped-to-support-secure-ai

3. (2025, December 10). OWASP Top 10 For Agentic Applications 2026. OWASP. https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/

ABOUT THE AUTHOR

Tech Mahindra

Ravi Sharma - Sr. Vice President & Global Head - Microsoft Business Feed | LinkedIn

Vinod Radhakrishnan - Head AI-Strategy and Growth - Vinod Radhakrishnan | LinkedIn

Microsoft

Nitin Sharma - Global Director, Microsoft Nitin Sharma | LinkedIn

Hari Krishnan T.M - Technical Director - Cloud Solution, Microsoft Hari Krishnan T.M | LinkedIn

BEST PRACTICE REPORT

Closing The AI Trust Gap With Service Providers

Lack Of Trust Is Fueling AI Concerns Between Clients And Their Services Partners

November 4, 2024

By Bill Martorelli, Christopher Condo with Linda Ivy-Rosser, Michele Goetz, Lauren Alexander, David Morrison

FORRESTER®

Summary

The need to explore and exploit generative AI (genAI) for competitive advantage and productivity gains brings with it new risks to the enterprise. These include emerging concerns about genAI's inputs and outputs. Global systems integrator (GSI) partners play a growing and important role in exploring and mitigating these risks as they deliver genAI-based solutions for their clients. GSIs can help enterprises cultivate a positive and trustworthy image of genAI and encourage its widespread adoption by understanding and communicating the risks of their genAI solutions. This report investigates how genAI can help you establish trust in the context of software developed by service partners.

GenAI Creates Opportunities And Risks For Enterprise/Service Partnerships

Excitement about genAI's potential to impact software development and operations has been spurred by expectations of significant productivity gains, as well as the tantalizing prospect of new revenue generation. Yet enterprises are torn between the tremendous promise of genAI and its risks. This is particularly true given that the governance dynamics and regulatory frameworks are still unfolding. The emerging risks are multifold, but two challenges stand out in the context of software developed by enterprises themselves or with systems integrators (SIs) with the help of genAI. These concern reliability and transparency related to genAI's inputs and outputs:

1. How can enterprise customers be sure that software products created for them by SIs with the help of genAI will act as intended and will minimize the threat of hallucinations or other unwanted outcomes?

2. How can enterprises be sure they are not facing potential legal risks to the business through potential misuse of proprietary intellectual property (IP), whether in materials used to train LLMs or in generated code or content?

As builders of solutions that encompass genAI, GSIs play a variety of roles in the emerging landscape, including as educators, implementers, and guarantors. Moreover, several large GSIs, including Accenture and Cognizant, have announced significant investments in genAI, which has emerged as an opportunity for growth for GSIs in a market that's otherwise constrained on the discretionary side. But in the longer term, GSIs are poised to play an invaluable role as enterprises continue the lengthy process of assimilation. Moreover, genAI will have significant implications on how GSIs organize themselves and deliver software to their clients.

GenAI's "Black Box" Problem Means A New Level Of Transparency Is Needed

GenAI may itself lack transparency in that its inner workings are opaque, making it difficult to understand how it derives decisions and if the underlying data or algorithms it uses are biased. GenAI builders and their customers need an expanded dialogue to address current challenges.

- **Customers want to make use of genAI, but some are still nervous.** The early days of genAI and open source are similar, with questions still unresolved. Enterprises are clearly tempted by the potential, but they're also fearful of its consequences and risk. In some cases, suppliers report, customers are placing "No AI" notices

in their RFPs! Accenture has reported $900 million in genAI-related revenues but acknowledges that most activity is still in the proof-of-concept stage. But the genie is out of the bottle, and enterprises' ability to resist genAI is potentially limited. "Engineers are engineers," observed one senior executive at a software engineering company. "They're going to do whatever they are going to do."

- **Ambitions to monetize genAI raise the stakes.** Significant impact from genAI on developer productivity and business process execution is widely anticipated. While enterprises expect productivity gains, they're anxious to monetize genAI's potential. This will likely entail implementing genAI outside the comfort of on-premises confines, making the risks more tangible. Moreover, the commercial model of genAI remains unclear. The possibility remains that capabilities of genAI solutions will be oversold, leading to high cost and underperformance.

- **GSIs and customers are not clear on how to assign risks.** How to address the risks of genAI is being discussed in many conversations that enterprise customers are having with their GSIs, but it's not clear yet how these risks should be assigned. The current situation also resembles the early days of the General Data Protection Regulation (GDPR), when the respective responsibilities of data producers and consumers were not well understood. Government initiatives like the EU AI Act will help clarify respective responsibilities, but for now a period of watching and waiting is in store. While a wholesale shift in work-for-hire contracts is not imminent, suppliers and their customers are watching for further clarification that will emerge from the courts and other deliberative bodies.

Enterprise Customers Are Proceeding Carefully

Enterprise customers are keenly aware of the risks associated with genAI. But they also understand that doing nothing is not an acceptable option. As a result, they are proceeding carefully. The ambitious plans of experimentation to production require guardrails, favorable term negotiations, and third-party mitigation strategies, including the following:

- **No training allowed.** A preference for isolated, on-premises implementations is clearly evident among nervous companies, which want nothing to do with hyperscaler versions. Others are willing to work with hyperscaler versions as long as there is strong indemnification or other guarantees. One consistent message from suppliers like NTT DATA is that customers don't want public models to be trained on their private data. Moreover, CEOs want access to success stories that have yet to arrive.

- **Relying on a "human in the loop" to protect against adverse outcomes.** GenAI's probabilistic behavior requires human reviewers; however, that human becomes accountable. This is the same as in traditional software development, but there are limits to what a single individual can do. Validation and verification at the line-of-code level is impractical; moreover, developers cannot simply comment on every line of code generated by LLMs. But what does due diligence look like in automated testing? The breadth of potential output is wider, and the relationship between inputs and outputs is less intuitive. "We're getting closer to an infinite number of monkeys on an infinite number of typewriters," said one supplier. "We are going to have collisions."
- **Red-teaming and adversarial testing "brute force" tactics.** One SI observed that when you don't know how to do something perfectly, the tendency will be to brute-force it with approaches like red-teaming, in which a number of individuals seek to purposely "break" a system, or adversarial testing. This can make a great deal of sense in certain circumstances. The U.S. White House's "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence" addresses the role of red-teaming explicitly.

Key Questions Being Raised By Enterprise Leaders And Their GSI Partners

GenAI poses a different set of problems than traditional software development. Much is the same, but some key things are different. The industry as a whole is currently sorting through these issues. How should these risks be apportioned? What does due diligence look like in the era of genAI? The very unpredictability of genAI is key to its promise — and peril. As a result, verification, validation, and overall governance strategies must be carefully calibrated. Enterprises must consider the following:

- **What will comprise meaningful test coverage in a genAI-infused dev environment?** Testing for deterministic systems is more straightforward than testing for probabilistic systems, and the relationship between inputs and outputs is less intuitive, while the breadth of potential output is wider.
- **Are there limits to what a "human in the loop" can do?** Many enterprises envision placing a human in the loop as a way to mitigate risks. But the ability of a single human is limited, whereas the potential risks of genAI are manifold. Implementers of genAI must take care to avoid a scenario in which a human in the loop will serve merely as a potential scapegoat should things go wrong.
- **How "heavy" should governance be?** Enterprises must take considerable care in formulating governance strategies for genAI. The risks would justify potentially

infinite control, but too much governance would likely stifle genAI's potential for innovation. Brillio, for example, explores the idea of "minimal viable governance" in its genAI engagements.

Systems Integrators Are Demonstrating Commitment And Rising To The Challenge

SIs are having many conversations with customers and going through careful deliberations about when to take responsibility, and when not to. Note the following trends:

- **Existing IP protections suffice — at the moment.** Service companies are finding that specific terms and conditions for genAI are not required for now, but that existing provisions, including those addressing the protection of IP and customer data privacy, are enough for the moment (although some customers will seek explicit guarantees that the model won't be trained on the customer's data). GSIs want to maintain the same level of indemnification risk as those held by the LLM providers, but not more. But the model is slowly being challenged, and existing provisions are being stretched to cover genAI explicitly. More genAI-specific terms and conditions are expected in the future.
- **Indemnification generally "flows through" from LLM providers — for now.** Enterprise customers working with GSIs do not need to establish indemnification with their service partners, at least not yet. Usually the providers of tools and LLMs provide indemnification, allowing the GSIs to lean on the protection provided by large suppliers like Microsoft that offer strong indemnification for customers. But the implications remain largely untested and may still be of concern to your chief security officer.
- **Suppliers help customers reduce risk with access to "safe" services and best practices.** GSIs like Softtek are helping their enterprise customers with ways to implement genAI solutions safely, by employing practices like the physical isolation of LLMs, retrieval augmented generation, separation of responsibilities for LLMs, using code filters for GitHub, establishing so-called cognitive firewalls to validate and verify genAI systems, and other ways to reduce risks. Another key role is reducing the risk of IP leakage. Cognizant, for example, helps enterprises build cognitive firewalls to monitor LLM inputs and outputs to help identify the risk of protected IP being employed inadvertently.
- **Risk/reward calculations will rely on the choice of commercial model.** The willingness of GSIs to accept risk is directly proportional to their possible gains.

Accordingly, selecting an appropriate commercial model is crucial. For example, suppliers are not likely to accept a lot of risk for staffing-based relationships, instead reserving that for more-managed engagements where shared risk and opportunities are present.

The Principles Of Shared Risk And Shared Responsibility Are Relevant

We're entering a period of shared risk and shared responsibility (see Figure 1). Assess your appetite and act accordingly. With experience with regulations like the GDPR, the industry has arrived at a good understanding of respective responsibilities between data providers and data processors. With genAI, there will similarly be a period of uncertainty as responsibilities between customers, deployers, and SIs become better understood. In the meantime, be prepared for a period of uncertainty between SIs, customers, and LLM providers and:

- **Engage your GSI partners in dialog, but don't expect them to have all the answers yet.** GSIs are looking for guidance as well and waiting for the fog to clear. In particular, they're looking at regulatory guidance like NIST AI 600-1 and genAI methodologies emerging from major vendors such as Google's Secure AI Framework (SAIF) to help light the way (see Figure 2). Approach your discussions with cocreation and partnership in mind.
- **Recognize that the hallucination problem will be difficult to solve.** Early genAI-infused systems have been prone to unpredictable or outright erroneous outputs. For example, the Stanford Institute for Human-Centered AI found in 2023 that general-purpose chatbots were hallucinating between 58% to 82% of the time in response to legal queries, including citing nonexistent cases in court papers. Presumably, this helped convince US Supreme Court Chief Justice John Roberts to warn of the risks of genAI in his 2023 Year-End Report on the Federal Judiciary. Although much can be done to minimize their occurrence, hallucinations may prove impossible to eliminate completely. Some suppliers have built tools that attempt to detect or prevent potential hallucinations. For example, Tech Mahindra's TechM VerifAI, which it sells as a software solution with services, is intended to prevent LLM "drift." That genAI systems won't experience hallucinations will be hard to guarantee and put into writing, and, as a result, sign-offs and transfer issues will be complicated.
- **Think internally before extending genAI to customer-facing systems.** N-iX recommends that its customers concentrate first on internal systems as opposed to external systems, observing that systems that include a lot of language translations are highly prone to hallucination. This dovetails nicely with the established

preference to pursue genAI in a protected, secure internal fashion.

- **Understand that genAI literacy is an important part of the story.** One impact of genAI will be a significant focus on ongoing education and training, along with organizational change management. “Lots and lots of training is kind of a given,” observed one GSI. You should anticipate the involvement of nontraditional subject-matter expert (SME) participants, including anthropologists and psychologists. Use your GSI partners to broaden these discussions.
- **Match governance with the needs of the workload.** Finding the right approach to governance will require fine calibration. Brute-forcing every workload having to do with genAI will thwart the very qualities that are drawing enterprises to genAI in the first place.
- **Recognize the limits of humans in the loop.** It is impractical to expect that a single individual can mitigate all of the risks associated with genAI systems by simply being “in the loop.” Most code reviews, for example, focus on the quality of the code itself, not its potential for IP leakage. Moreover, a single person cannot be responsible for documenting and commenting on all AI-generated code.

Figure 1

Suppliers, GSIs, And Enterprises Own Different Degrees Of Shared Risk

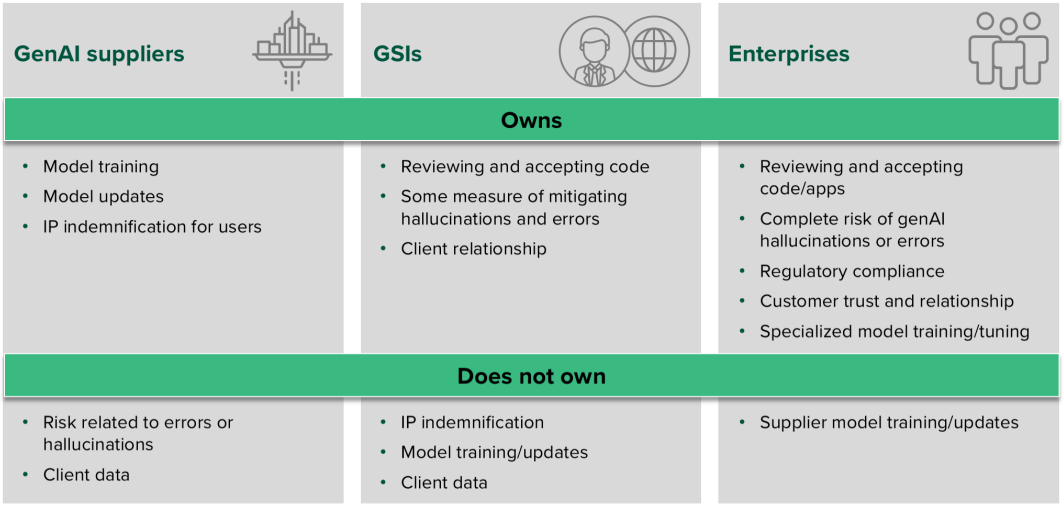

Figure 2

Standards And Regulations Pertaining To GenAI

Issuing body	Name of initiative	Comments
The White House	Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence	Issued October 30, 2023, the more than 100-page executive order provides new ways to report harmful AI practices and hold irresponsible managers of AI systems accountable. However, it is only relevant to a subset of institutions.
European Union	EU AI Act	Effective August 1, 2024, the EU AI Act establishes a common regulatory and legal framework for AI within the EU. The act is viewed as a somewhat lightweight framework aimed at the largest set of risks.
U.S. Department of State	Risk Management Framework (RMF) for Artificial Intelligence and Human Rights	Issued on July 25, 2024, the RMF is intended as a practical guide to help organizations design, develop, deploy, use, and govern AI in a manner consistent with respect for international human rights.
U.S. National Institute for Standards and Technology	NIST AI 600-1: Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile	Sent out for comment in April 2024 and issued in response to the White House executive order, NIST AI 600-1 is a comprehensive risk management framework profile intended to complement the U.S. State Department's RMF. Observers say the model appears comprehensive in nature.
MITRE	ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems	Begun in 2020, ATLAS provides a knowledge base sharable across government, industry, academic, and international allies to anticipate and address national AI challenges. As such, it primarily focuses on security. ATLAS is modeled on MITRE's ATT&CK framework and its tactics, techniques, and procedures (TTPs).
Google	Secure AI Framework (SAIF)	Introduced in June 2023, SAIF is designed to address concerns for security professionals, such as AI/ML model risk management, security, and privacy. It is viewed as an essentially bottom-up approach that's friendly to genAI developers.
Cyberspace Affairs Commission of China (CAC)	Interim Measures for the Management of Generative Artificial Intelligence Services	Effective August 2024, the interim measures establish a set of rules to be followed by anyone who develops or uses genAI products to deliver services to Chinese consumers. It is viewed as particularly comprehensive, but strict.

What GenAI Means For The Coming Impact On SI Operating Models

GenAI will impact both customers and SIs, as well as the way they interact. Its impact will include more productivity, more code being delivered, and more code being ingested to train the models. Note the following trends:

- **GenAI will impact team-based delivery models.** GenAI will impact the way enterprises and GSIs organize teams. By helping them automate and streamline intrateam processes, genAI will contribute to overall team cohesion and reduce the need to colocate teams for the same purpose. This will make it easier for enterprises and GSIs to establish teams consisting of far-flung resources.

- **GenAI will shift to a more SaaS-based delivery model.** The classic approach to "work for hire" has entailed significant and lengthy handover processes, including end-user acceptance testing and other criteria. With genAI, however, work products will be harder and harder to simply hand over. Work-for-hire trends have already been impacted by DevOps concepts combining build and run. Probabilistic genAI systems take it one step further, making custom-built software more like software as a service (SaaS), in that it won't be turned over in transition but be maintained permanently by the supplier.
- **GenAI will alter staffing pyramids of GSIs.** Historically, GSIs like Accenture have relied on a classic pyramidal model, in which a few staff at the very top are supported by hordes of resources at lower levels, typically supported by numerous low-skilled resources at the very bottom. With genAI, observes Cognizant, many of the basic manual tasks at the bottom will likely be automated, while those at the top of the pyramid will have greater reach by virtue of genAI's multiplicative power. This will have lasting implications on the operations of GSIs and their interactions with their customers.

Supplemental Material

Companies We Interviewed For This Report

We would like to thank the individuals from the following companies who generously gave their time during the research for this report.

Accenture

Brillio

Cognizant

EPAM

N-iX

NTT DATA

Softtek

Tech Mahindra

FORRESTER®

We help business and technology leaders use customer obsession to accelerate growth.

FORRESTER.COM

Obsessed With Customer Obsession

At Forrester, customer obsession is at the core of everything we do. We're on your side and by your side to help you become more customer obsessed.

Research

Accelerate your impact on the market with a proven path to growth.

- Customer and market dynamics
- Curated tools and frameworks
- Objective advice
- Hands-on guidance

Learn more.

Consulting

Implement modern strategies that align and empower teams.

- In-depth strategic projects
- Webinars, speeches, and workshops
- Custom content

Learn more.

Events

Develop fresh perspectives, draw inspiration from leaders, and network with peers.

- Thought leadership, frameworks, and models
- One-on-ones with peers and analysts
- In-person and virtual experiences

Learn more.

Contact Us

Contact Forrester at www.forrester.com/contactus. For information on hard-copy or electronic reprints, please contact your Account Team or reprints@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research, Inc., 60 Acorn Park Drive, Cambridge, MA 02140 USA
Tel: +1 617-613-6000 | Fax: +1 617-613-5000 | forrester.com

ABOUT TECH MAHINDRA

Tech Mahindra (NSE: TECHM) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 150,000+ professionals across 90+ countries helping 1100+ clients, TechM provides a full spectrum of services, including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal in recognition of actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federations of companies.

For more information on how TechM can partner with you to meet your scale at speed imperatives, please visit Tech Mahindra | Scale at Speed