

The agentic AI shift: What leaders must prepare for in the next 12–24 months

AI agents are the new attack surface. Secure your agents, the data they touch, and the actions they take.

Just 21% of organizations report achieving AI maturity—where AI is fully deployed across cybersecurity, risks are assessed, effectiveness is measured via KPIs, and executives have regular visibility into AI’s impact on preventing and reducing cyberattacks.¹

¹ Ponemon Institute, *Managing Risks and Optimizing the Value of AI, GenAI & AgenticAI*, February 2026. Sponsored by OpenText.

Agentic AI represents a fundamental shift in how organizations use artificial intelligence—from systems that assist humans to systems that autonomously gather information, make decisions, and act. Over the next 12 to 24 months, these systems will rapidly move from experimentation to operational reality across IT, security, digital experience, and business operations.

This shift promises dramatic gains in speed, efficiency, and scalability. But it also introduces new categories of enterprise risk that existing governance, identity, and security models were not designed to handle.

Leaders who prepare now will unlock productivity without sacrificing control. Those who don’t risk autonomous systems acting faster than their ability to detect, intervene, or recover.

What is agentic AI—and is it different?

Traditional AI systems analyze, recommend, or generate outputs for humans to approve. Agentic AI systems operate with intent.

They:

- Gather information across systems
- Make decisions aligned to defined goals
- Act without human approval

Autonomy changes accountability—agentic AI behaves more like a digital employee than a traditional tool. Decisions are probabilistic, not deterministic, making outcomes harder to fully predict or pre-approve. Security models designed for static software and human users no longer apply cleanly. In effect, agentic AI systems behave like digital actors inside the enterprise, not just tools. This distinction matters because autonomy changes the risk profile entirely.

Where agentic AI will appear first

Agentic AI adoption will not be evenly distributed. Early “landing zones” are emerging where autonomy delivers immediate value:

- **IT operations:** Automated remediation, infrastructure optimization
- **Security operations:** Triage, response orchestration, threat containment
- **Digital experience and content ops:** Personalization, workflow automation
- **Supply chain and logistics:** Demand balancing, exception handling
- **Data and analytics:** Continuous insight generation and reporting

Early adoption areas combine speed, access, and autonomy, and agentic AI can act faster than any human team, amplifying the impact of any error or misuse. These environments often involve sensitive data and production systems, increasing regulatory and operational exposure.

The agentic AI revolution has begun



Definition

Goal-driven systems that can gather information, make decisions, and then act.



Landing zones

Anywhere. Content ops, IT ops, security ops, customer and digital experiences, supply chains, product development, data management analysis and reporting, and more.



Highly productive

Dynamic and sometimes highly privileged access to control and change data, systems, processes, policies, and even devices at the speed of AI to achieve goals.

The productivity upside is real—and significant

Organizations deploying agentic AI report:

- Faster decision cycles
- Reduced operational friction
- Improved consistency and scale
- Lower reliance on human intervention for low-risk tasks

In isolation, these gains look transformative. The challenge is what happens when autonomy intersects with access, data, and control.

Why agentic AI expands the attack surface

Agentic AI doesn't just consume data—it touches systems.

Key risk shifts include:

- **AI attack surface expansion:** Models and data pipelines introduce new pathways for data leakage, manipulation, and control bypass
- **Access drift:** Agents may change permissions dynamically to achieve goals
- **Unintended consequences:** Literal goal execution vs. business intent
- **Privilege accumulation:** Agents inheriting or escalating access over time
- **Speed of impact:** Errors or misuse propagate faster than manual response

In real-world scenarios, an agent optimizing for efficiency can unintentionally disrupt critical services, expose sensitive data, or override established safeguards.

Why traditional controls fall short

Most enterprise controls assume:

- Humans initiate actions
- Access is relatively static
- Decisions are reviewable before execution

Agentic AI breaks these assumptions. Controls designed for users and apps struggle when faced with autonomous, decision-making entities operating continuously and at machine speed.

What leaders must prepare for now

Over the next 12 to 24 months, organizations must evolve from “AI enablement” to AI governance at machine scale.

Key preparation areas:

- **AI AppSec preparedness:** Detect and control AI-specific vulnerabilities
- **Identity for non-human actors:** Treat AI agents as identities, not tools
- **Dynamic access controls:** Continuously evaluate what agents can do
- **Data-centric security:** Control not just access, but data usage and movement
- **Visibility and auditability:** Understand agent decisions and actions after the fact
- **Containment and recovery:** Assume failure and design for rapid intervention

This is not about slowing innovation, it’s about making autonomy safe at scale.

Trust by design: Build the runway first controls to allow for secure speed



Identify sensitive data

Know where your important information lives and protect it.



Secure your apps

Prevent vulnerabilities in agent and tool code.



Secure identities

Make sure people and systems only have the access they need.



Be ready to respond

Have a plan to act fast if something goes wrong.



Monitor behaviour

Keep an eye on what’s happening, so nothing slips through.

6 strategic questions every executive should ask

1. Are we building agents securely?
2. Where will agentic AI operate in our environment first?
3. What systems could it access—directly or indirectly?
4. How do we detect unintended behavior in real time?
5. Who is accountable when autonomous systems make decisions?
6. Can we stop, constrain, or roll back an agent instantly?

If these questions don’t have clear answers, preparation has not yet begun. Download the [Secure AI readiness checklist](#) to assess your enterprise’s readiness.