

# Control and manage your *agentic* *workplace*

# The agentic workplace: *where belief meets adventure*

Many organizations have bold ambitions for delivering business value with AI.

They want to embed Gen AI into core processes to increase operational efficiency and productivity at scale. Some plan to use AI to make it easier for teams to deliver more tailored and responsive customer experiences. Many also expect it to unlock entirely new ways of operating, propelling innovation to new heights, and redefining how value is delivered in their sector.

And in the workplace, pioneering organizations expect greater automation, streamlined workflows, and faster insights to make collaboration simple and empower everyone to focus on meaningful work.

To accomplish these goals, they're deploying AI agents to augment human capabilities and act autonomously to complete business processes. But as the agentic ecosystem expands, it can also increase the risk of exposure. The agentic workplace will only generate business value if it has the right governance foundation.

By the end of 2026, **40%** of enterprise applications will embed AI agents, up from less than 5% in 2025<sup>1</sup>

Gartner

## Without governance, more agents lead to more chaos

In many cases, organizations have rushed to deploy agents, but lack central visibility into their agentic workforce. They have no idea how many agents there are, who owns them, or what they're doing. Agents often have inconsistent or non-existent security measures, and the growing problem of shadow AI creates even more security and compliance risks.

In ungoverned environments, these issues multiply with each new agent. Ad hoc development, unclear ownership, and patchwork support mean agents struggle to deliver meaningful value or worse; they simply increase complexity and risk.

Aware of these risks, some organizations are delaying agent deployment until they know they can overcome the security and compliance challenges. But every day the agentic transformation waits is another day of falling behind competitors.

Other organizations sit somewhere between these two extremes. They're experimenting with agents in sandboxes, but these prototypes have no real impact on the business.

In this guide, we'll explore how to gain visibility and control of your AI ecosystem to create a unified, secure, and well-governed agentic workplace that delivers value without increasing risk.

<sup>1</sup><https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>



# What is the *agentic workplace*?

In the agentic workplace, AI is at the core of the digital environment, seamlessly connecting all applications and processes into a unified ecosystem. People access this ecosystem through a simple interface, with AI agents working behind the scenes to retrieve information, complete tasks, or even autonomously orchestrate entire business processes.

These multi-agent environments bring together numerous AI capabilities to streamline workflows and simplify everyone's working day.

For example, IT helpdesk agents can connect platforms like ServiceNow and Microsoft 365, answering questions with custom knowledge and resolving tickets quickly to keep users productive and free skilled engineers to focus on the toughest issues. HR agents can unite Microsoft 365 with systems such as Workday to guide new hires seamlessly through onboarding, so they can hit the ground running. And sales agents can

rapidly analyze and report on vast quantities of data in CRM platforms like Salesforce to provide timely, targeted insights that help salespeople focus on the most promising leads.

AI agents surface the right information instantly, take on routine and repetitive tasks, and proactively remove friction from the working day. Managers benefit from agents that continuously scan trends and anomalies, helping them make faster, more accurate decisions and stay focused on the KPIs that drive performance. And leaders gain immediate, data-rich answers to critical questions without searching through dashboards or reports, allowing them to steer the organization with clarity and confidence.

To sustain these benefits at scale, however, the agentic workplace must be underpinned by strong governance, a secure platform, solid data foundations, and a mature operating model. Organizations must manage shadow



AI and agent sprawl, ensure endpoint readiness, embed compliance and identity-centric security, and adopt evergreen operations and CoE guardrails.

With these foundations in place, agents can operate safely, reliably, and consistently, enabling the agentic workplace to deliver impact for everyone in the organization.

## Governing the *agentic workplace*

The management of the agentic workplace spans several areas, including agent DevOps, business and employee enablement, business value reporting and agentic governance and security.

As the drive to embed agents in day-to-day operations becomes a strategic imperative for staying competitive, strong governance is becoming a strategic priority. Every application, workflow, and agent must operate within a centrally governed platform to ensure security, observability, and alignment with organizational policy.

With an integrated governance framework, organizations can turn the chaos of agent sprawl into a secure, observable, and continuously improving agentic workplace. This is where solutions like Microsoft Agent 365 are invaluable, providing a single control plane to observe, manage, secure, and assess every agent in the organization. A unified control plane enables IT teams to ensure agents are safe and compliant, and connect them directly to measurable productivity gains, user experience improvements, and strategic business outcomes.

### From chaos to control

A strong governance framework and tools like Agent 365 that are embedded in the Microsoft security ecosystem alongside familiar tools such as Entra ID, Purview, and Defender, enable organizations to:

- **Control every agent at scale**, with a single pane of glass to govern and manage all Microsoft, third-party, and custom-built agents.
- **Secure every prompt**, with a unified defense layer ensuring agents operate securely within defined

boundaries, and preventing unauthorized agents from bypassing compliance controls.

- **Standardize agent deployment**, linking design, operation, and continuous improvement to business goals and governance requirements.
- **Manage the entire agent lifecycle**, establishing clear ownership, support accountability, and continuous improvement at scale

# The high-performance agentic workplace: *unified, secured, controlled*

In the agentic workplace, everyone from the boardroom to the frontline benefits from AI agents augmenting their day-to-day work, automating routine tasks, and accelerating decision-making. But as more people gain access to Gen AI, organizations must ensure agents operate within strong guardrails.

Policy enforcement across all agents, managing identity and access, detecting and mitigating threats, and protecting sensitive data, organizations can confidently scale agent adoption without compromising security or compliance.

When everyone can safely use AI agents supported by enterprise-grade guardrails, continuous monitoring, and built-in protections, deployment confidence increases, adoption accelerates, and business impact multiplies.


**Regain control** with complete visibility and centralized command of the agentic workforce. Inventory, monitor, and manage every agent with instant insights from a single dashboard instead of maintaining manual registries.

**Build trust** by ensuring reliable operations and agent accountability. Apply structure and consistency to the

agent lifecycle, make every agent accountable to an owner, and address issues in hours, not weeks.

**Scale with** confidence by delivering security and compliance by design. Ensure no agent can operate outside enterprise guardrails, every AI decision is auditable, and sensitive data is demonstrably protected.

**Boost business** agility by deploying new AI initiatives in weeks, with a governance framework already in place to accelerate deployment and adoption without increasing risk.



## Critical success factors for the *agentic workplace journey*

Creating an agent-enabled workplace is more than just deploying AI. It's about embedding governance, lifecycle management, and security into every layer of the digital environment. This enables organizations to evolve from deploying isolated AI tools to enable a fully governed agentic workplace, where agents operate within a secure, compliant, and observable framework.

Success on this journey depends on aligning people, process, and technology, anchored by centralized control, responsible AI, and identity-centric safeguards. The most important actions to take include:



### People

- Drive adoption by understanding user needs and equipping them with the skills to work safely and effectively with AI agents.
- Establish clear ownership through agent identity and role-based access.
- Build trust through transparency, emphasizing responsible AI principles and positioning agents as governed collaborators that enhance, not replace, human work.



### Process

- Focus on high-value use cases and define KPIs that measure both business impact and compliance.
- Use a centralized governance model, such as a Center of Excellence, to manage agent lifecycles, enforce policy, and ensure responsible deployment.
- Embed observability, feedback loops, and red teaming into operational processes to continuously improve agent performance and trustworthiness.



### Technology

- Establish a secure foundation with a centralized control plane, such as Agent 365, to manage agent identity, access, and lifecycle.
- Enforce least-privilege access, DLP, and encryption, and integrate tools such as Defender, Purview, and Sentinel for real-time monitoring, threat detection, and compliance.
- Embed governance by design, ensuring agents are safe, compliant, and aligned with enterprise policies from day one.

# Agentic governance *in action*

With Capgemini’s proven approach to agentic governance, one of Australia’s largest superannuation funds scaled Microsoft 365 Copilot to 1,200 employees across six locations.



## The Challenge

Digital fluency gaps in the workforce meant there was inconsistent AI readiness across teams. Plus, without a governance framework for AI tools, the company faced heightened risks of data security breaches, shadow AI use, and uncontrolled adoption. Leadership needed confidence that scaling Copilot wouldn’t compromise compliance or overwhelm employees.



## The Solution

Capgemini delivered a structured enablement and governance program, including:

- Dedicated hub with 15 micro-training videos
- Network of Copilot champions across teams
- Governance framework for responsible AI usage
- Structured PoC for agent development

We also built three custom agents: a service desk agent for IT support triage, a risk assessment agent for compliance reviews, and a test case generation agent for QA automation.



## The Results

The service desk agent has measurably reduced the pressure on first-line support. The risk assessment agent now handles 70% of routine approvals autonomously. And the test case agent has saved over 2,000 hours of manual QA effort.

With Capgemini’s governance framework in place, the company has moved from cautious experimentation to confident, enterprise-wide AI adoption, all within a secure, auditable environment.

- **92%** active user adoption
- **~900** weekly Copilot users
- **92K** prompts per 4-week period
- **2,000+** hours saved (test case agent)

# The ideal starting point for *your agentic workplace*

The fastest way to build a connected agent ecosystem is to start with the platforms your organization already trusts. By using existing identity, data, and collaboration foundations, you ensure agents operate in a secure, managed environment from day one.

A centralized governance platform such as Agent 365 becomes the entry point, providing a unified registry of all agents, consistent security, access controls, and seamless integration.

With hundreds of enterprise connectors available, agents can securely access work data and systems, delivering answers and actions through a single, intuitive interface. When an employee asks a question, the control plane orchestrates the right agents behind the scenes.

This approach also prepares you for scalable agent development. No-code and low-code builder tools enable the rapid creation of simple agents, while pro-code frameworks support more complex, autonomous use cases. Both work under the same governance model, enabling quick wins and long-term evolution without switching ecosystems.

## The Capgemini difference

As an **official Microsoft launch partner for Agent 365**, Capgemini benefits from early access to Microsoft’s latest AI innovations. Our teams also co-develop reference architectures, governance frameworks, and deployment accelerators aligned with Microsoft’s roadmap, enabling organizations to scale AI agents securely and confidently.

Our strategic partnership with Microsoft extends to our co-created Agentic Industry Studio. It combines

Capgemini’s deep industry expertise and proven AI transformation playbooks with Microsoft’s AI platforms to accelerate value delivery in the agentic workplace.

In addition, Capgemini RAISE™ provides an enterprise-grade foundation and accelerators for building, integrating, and operating agentic AI at scale. This foundation enables organizations to confidently develop, orchestrate, and manage AI agents at scale, unlocking operational efficiency, integrated workflows, and connected experiences and accelerating the journey toward an AI-native, multi-agent workplace.

This is where belief meets adventure: bold ambition underpinned by the discipline and foundations required to make the agentic workplace real.

## Begin your adventure

Capgemini helps you establish this foundation by setting up the governance control plane, integrating identity and security, and building a Center of Excellence for agent lifecycle management.

We then focus on high-value use cases and guide your teams in blending low-code and pro-code development to scale responsibly and confidently. Our experience and frameworks accelerate your journey toward an agentic workplace that empowers every employee while maintaining the right guardrails.

To learn more about how we can help you build and govern your agentic workplace, [get in touch with our experts](#).

## Author



**Lukasz Ulaniuk**

Global Portfolio Director,  
Workplace & Employee Experience,  
Cloud Infrastructure Services

For more details, contact:

*[Infra.global@capgemini.com](mailto:Infra.global@capgemini.com)*

## About Capgemini

Capgemini is the business transformation partner for enterprises in the age of AI. We help organisations imagine and build an intelligent, sustainable future, combining AI, technology and human ingenuity to transform how they operate, innovate and grow. With unique end-to-end capabilities spanning strategy, technology, engineering and intelligent operations, we bring together deep industry expertise and market-leading capabilities in AI, cloud and data to turn ambition into measurable business outcomes at scale. Supported by a robust ecosystem of partners and nearly 60 years of expertise, Capgemini is a responsible and diverse global organization of over 410,000 team members in more than 50 countries. The Group reported 2025 revenues of €22.5 billion.

**Make it real.**

[www.capgemini.com](http://www.capgemini.com)

