Agentic Control and Management

Governing, Operating, and Securing AI Workforce.

Capgemini | Microsoft

The emergence of the digital workforce

AI agents growing exponentially - governance is now a board-level strategic imperative

Personal assistants that help individual users' complete tasks, E.G. A sales rep's AI helper or personal copilot. Scope: single user, low autonomy.

Human-in-the-Loop Agents that collaborate with employees on complex tasks or decisions, augmenting human judgment. Require human approval at key steps.

Autonomous Agents that execute end-to-end processes independently, with minimal human intervention. Highest governance priority.

> If agents become your Digital Workforce, who is your Digital Workforce Department?

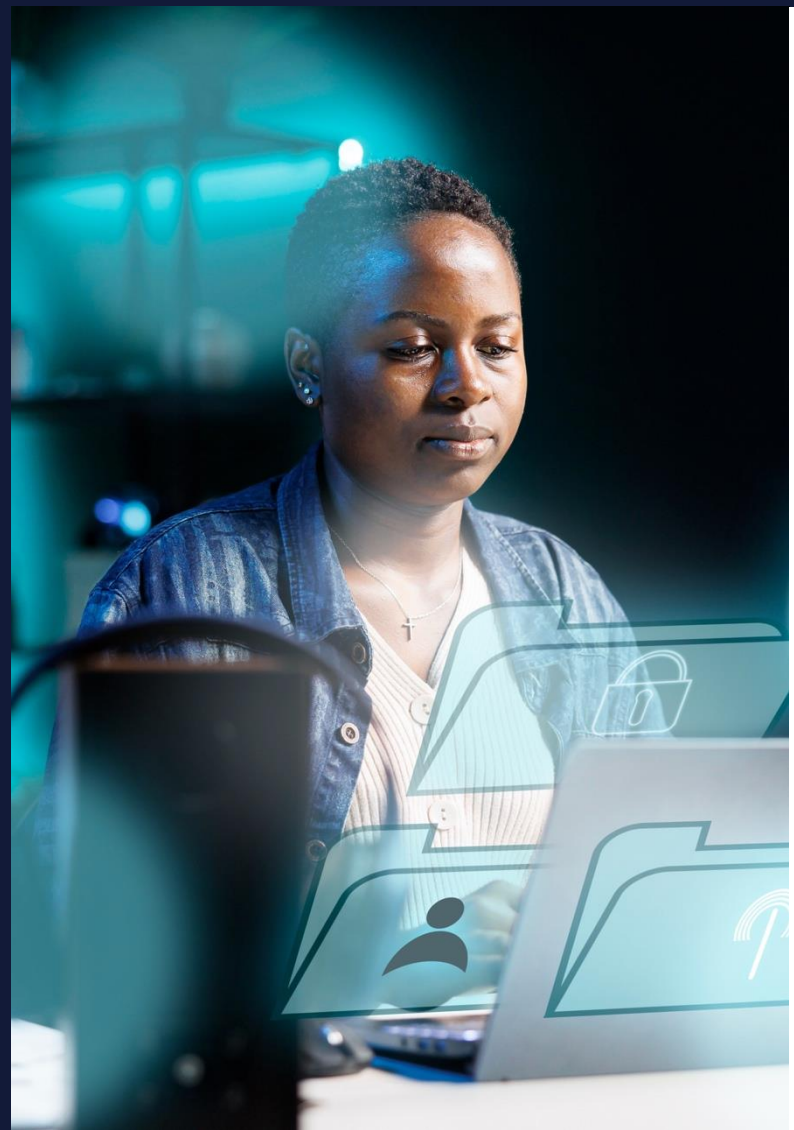

By the end of 2026, 40% of enterprise applications will embed AI agents, up from less than 5% in 2025

(Gartner press release, Aug 26, 2025)

Complex AI agent ecosystems increase enterprise risk

Can IT discover and manage agent proliferation for impact?

Are agents behaving correctly and appropriately within the enterprise?

Who / what are agents sharing sensitive information with?

Are the agents well governed and audited - what are my costs?

The full scope of agentic governance

A unified framework spanning strategy, security, operations, enablement, and infrastructure

Business strategy

The strategic foundation for unlocking new Agentic AI. Enables the design and execution of business models, innovation acceleration, and competitive advantage. This pillar aligns AI initiatives with P&L outcomes, market positioning, and long-term growth levers.

Outcome: AI becomes a core driver of business transformation, delivering measurable value across revenue, cost, and innovation metrics.

AI transformation strategy

The blueprint for scaling agentic AI through structured change and operational readiness. Integrates change management, operating model evolution, and workforce transformation to embed AI into the fabric of the enterprise

Outcome: organizations achieve sustainable adoption, empowered teams, and agile delivery models that accelerate time-to-value and ensure long-term success with agentic AI.

Agentic Governance & Security

The control plane, security, and compliance layer for all AI Agents enterprise-wide. Registry, Blueprints, Identity Management, DLP, threat detection, Responsible AI.

Outcome: every agent operates safely, transparently, and effectively and is 24x7 monitored for resilience

Agentic Devops

Manages the full agent development: conception through retirement. Agent Factory approach: structured development using CI/CD pipelines. Continuous monitoring, keeping agents current and performant.

Outcome: operational DevOps for AI agents.

Business & Employee Enablement

Drives adoption and effective human-AI collaboration. Training programs, AI academies, change management, role redefinition, and employee playbooks.

Outcome: employees augment work with AI naturally and confidently.

Business Value Reporting

KPI frameworks, executive dashboards,
ROI analysis per agent. Identifies highest-return agents; flags underperformers for retirement.

Outcome: revenue, cost savings, NPS, time-to-decision. Governance drives value, not just compliance.

Raise AI platform

The foundational infrastructure and engineering framework for building, deploying, and scaling reliable AI solutions across the enterprise. Modular architecture, reusable accelerators, MLOps pipelines, data governance, and model validation tooling.

Outcome: AI systems are robust, scalable, and aligned with enterprise-grade performance, security, and compliance standards.

Data

The unified data framework underpins all agentic capabilities. It harmonizes ingestion, integration, quality, governance, lineage, and semantic consistency across structured and unstructured sources: AI-ready by default

Outcome: agents operate on trusted, governed, data-enabling accurate reasoning, compliance by design, and scalable AI-driven insights.

Agentic governance takes center stage in AI investment

60% of organizations are already exploring **Agentic AI systems. Governance** is now a strategic imperative.

(Source: Capgemini research institute)

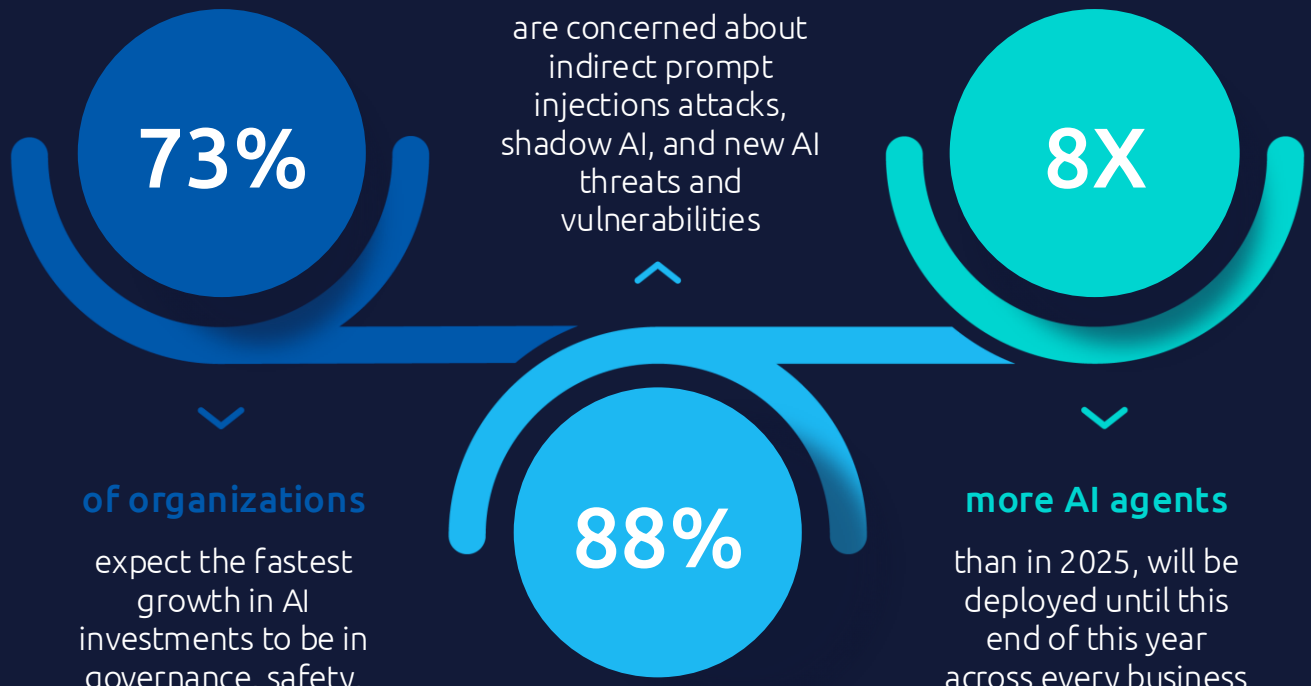

Capgemini's agentic governance directly addresses this urgency, enabling AI agents to scale safely alongside deployment.

2026: "The Year of Infrastructure Scale-Up" - global survey of 1,500 senior executives confirms accelerating AI investment in governance, data foundations, and workforce upskilling (Source: Capgemini Research Institute Global Executive Survey, 2025 – 2026)

Fragmented reality of today, full control of tomorrow

Without governance, more agents means more chaos - four critical pain points organisations face today

Today

Lack of visibility on agentic workforce

No central inventory or “control tower” - leadership doesn’t know how many agents exist, what they’re doing, or who owns them. Siloed deployments lead to duplicate efforts and zero enterprise-wide visibility.

High risk Agent's usage and shadow AI

Agents operate with inconsistent or nonexistent security measures. No DLP policies, no audit logs, high risk of data leaks and shadow AI. Employees spin up unauthorized bots using public AI services, bypassing compliance controls entirely.

Ad-hoc development and uncontrolled build quality

No standard agent blueprints or guidelines, agents built ad-hoc with wildly varying quality. No common lifecycle (design, deploy, monitor) - agents run on autopilot with no improvement mechanism until failure.

Scaling chaos

As agent count grows from 5 to 50+, problems multiply: overlapping functionality, unclear ownership, no support process. Without governance, more agents = more headaches, not more value.

Tomorrow

Unified control plane

Single pane of glass to inventory, govern, and monitor all agents, Microsoft, partner, and custom. Provides ownership, status, configuration, activity, and health visibility, enabling consistent, at-scale management.

Security with every prompt

Real-time threat detection, adaptive access safeguards, and shadow agent discovery into a unified defense layer. Agents operate securely within defined boundaries, respond to emerging risks instantly, and prevent unauthorized or unmanaged agents from bypassing enterprise.

Agents' deployment standards

Standardizes how agents are designed, deployed, and improved-linking business goals with governance from day one. Supports continuous optimization through telemetry, feedback loops, and structured updates, ensuring agents stay aligned with evolving needs.

Operational maturity

Manage lifecycle to design, deploy, monitor, and evolve AI agents at scale. Establishes clear ownership, support accountability, and continuous improvement across 300+ agents. Enables consistent performance, rapid issue resolution, and measurable value delivery.

From Ad-Hoc Beginnings to a scaled Agentic governance

A maturity journey from experimentation to enterprise-wide, governed AI agents

	Crawl	Walk	Run
Agent oversight	Few Agents under Governance, largely initiated by individual teams without central oversight	Initial inventory of agents established with majority of Agents having assigned owners, but monitoring is partial and fragmented.	Full visibility and centralized oversight via a unified control plane. Most of Agents are registered, monitored, and owned.
Process maturity & standards	No standardized development or deployment practices. Agent lifecycle is undefined and inconsistent.	Early governance processes introduced. Some lifecycle stages (e.g., deployment, monitoring) are defined but not enforced.	Enterprise-wide framework governs agent lifecycle from design to retirement. Continuous improvement is embedded.
Security & compliance	Security and compliance are reactive. No unified policies or controls across agents.	Basic controls applied to select agents. Compliance is ad hoc and auditability is limited.	Security, identity, and compliance controls are embedded by design. Full auditability and regulatory alignment achieved.
	Maturity: Level 1–2	Maturity: Level 2–3	Maturity: Level 4+

Accelerate your journey to secure, controlled, and scalable AI ecosystem.

A unified platform, embedded security, and continuous intelligence - powered by the existing Microsoft technology stack

01 Unified control plane

Capgemini sets up Microsoft Agent 365 as your single command center for every AI agent in the enterprise. All agents: from HR bots to finance copilots are registered, tracked, and managed through one console. Microsoft Foundry complements as the developer-facing control plane, providing observability, tracing, and runtime controls across the build-to-production lifecycle. M365 Admin Center provides centralized policy enforcement. No more shadow AI: every agent is visible, every deployment controlled, every update orchestrated.

02 Embedded security & compliance

We embed enterprise-grade control at every layer: from requiring Entra ID identities for agents, to enforcing DLP on all agent communications, to real-time SOC monitoring and vulnerability management with tools like Microsoft Defender. Any shadow bots are identified for onboarding or decommissioning. The result: no blind spots. You can pursue AI adoption at your own pace knowing that robust guardrails are in place (and operationally working 24x7) to find and contain breaches and ensure compliance.

03 Continuous monitoring & improvement

Capgemini establishes structured processes for agent owners to continuously monitor performance, collect user feedback, and trigger refinements: reconfigurations, prompt updates, or retirement. When an agent falls short, alerts fire and a defined improvement cycle kicks in. Your AI agents are never static: they evolve within a governed framework that ensures every change is tracked, tested, and approved. Regular evals assess agent performance, ensuring every agent delivers measurable value.

Together, these three solutions form a single, integrated governance framework turning ungoverned agent sprawl into a secure, observable, and continuously improving AI operations.

Core benefits of agentic governance - unified, secure, in control

Four core benefits - delivered across platform, security, and operations dimensions

01 Complete visibility & centralized command

Every agent, from HR FAQ bots to finance Copilots. Policy changes applied to all agents at once, admin overhead doesn't grow with agent count. Audits and "what AI do we have?" reporting become trivial.

02 Secure & compliant by design

No agent can operate outside guardrails: DLP, identity control, and monitoring ingrained from day one. Regulators and internal audit satisfied - every AI decision logged, sensitive data demonstrably protected. Zero critical AI incidents achieved by clients post-implementation.

03 Reliable operations & agent accountability

Structured lifecycle process (design → deploy → monitor → improve) ensures consistent agent quality and reliability. Every agent accountable to an owner, issues addressed immediately with a clear playbook. User trust increases: employees and customers experience agents as dependable and well-supported.

04 Business agility & future-proofing

Governance framework already in place, new AI initiatives deploy in weeks instead of months. Framework aligned to Microsoft's roadmap: new Copilot capabilities and autonomous agent features adopted seamlessly. Confidence in governance drives adoption; company can be an AI-first mover in its industry.

Real-world impact - AI agentic governance in action

How one of Australia’s largest superannuation funds scaled M365 Copilot to 1,200 employees with Capgemini’s governance approa ch

Problem

The challenge

- One of Australia's largest superannuation funds needed to roll out Microsoft 365 copilot to 1,200 employees across 6 locations.
- The workforce faced a significant digital fluency gap, with inconsistent ai readiness across teams.
- No governance framework existed for ai tools, creating risks around data security, shadow ai usage, and uncontrolled adoption. Leadership needed confidence that scaling copilot wouldn’t compromise compliance or overwhelm employees.

Solution

Capgemini’s approach

- Capgemini delivered a structured enablement and governance program:

1. A dedicated hub with 15 micro-training videos, copilot champion network across teams
2. Governance framework for responsible AI usage
3. Structured proof-of-concept for agent development.

- Three custom agents were built: a service desk agent for it support triage, a risk assessment agent for compliance reviews, and a test case generation agent for QA automation.

Outcome

Measurable results

- The risk assessment agent now handles 70% of routine approvals autonomously.
- The test case agent has saved 2,000+ hours of manual QA effort. The service desk agent measurably reduced first-line support pressure.
- All operating within a governed, auditable framework that gave leadership full confidence to scale further.

92% Active user adoption

~900 Weekly copilot users

92K Prompts per 4-week period

2,000+ Hours saved

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

The three pillars of agentic governance

Transformation & Configuration

P

- Centralized control plane - Enterprise-wide registry and governance via Agent 365, complemented by the Microsoft Foundry Control Plane for developer-level observability, tracing, and runtime insights across all agents (Microsoft, third-party, custom).
- Agent Blueprints: Ownership, Identity & Access Management - Every agent is deployed with a standardized Entra ID–based identity, assigned a Business Owner and IT Sponsor, and governed by defined permissions and access boundaries. Access is enforced via Agent 365 and validated through Microsoft Foundry - least privilege by design, full auditability throughout.
- Security Templates - Predefined policies for data access, DLP, logging, monitoring, and acceptable use: applied to every agent before and during deployment. Enforced continuously at runtime and updated centrally to adapt to evolving threats and regulations.
- Platform observability - Unified observability across Agent 365 and the Microsoft Foundry Control Plane, combining enterprise visibility (inventory, compliance, policy) with deep technical telemetry, tracing, evaluations, and runtime behavior insights.

ZeroTrust framework for AI

S

- Data protection & DLP - Rules enforcement on all agent interactions, sensitive data masked or blocked before output; retention policies enforced; Purview Sensitivity labels and data classification honored across agents. Data lineage for Agents decisions.
- Threat detection & runtime defense - Microsoft Defender, Sentinel SIEM, and custom Agent 365 alerts monitoring agent behavior, also with Agent 365 Runtime Shield. OWASP Agentic AI Risks Mapped to Agent 365 Detection Capabilities.
- Regulatory & responsible AI - Design-time Responsible AI assessments (bias, fairness, transparency); runtime audit logs for all agent decisions, aligned to GDPR, sector-specific regulations
- Conditional access & identity safeguards - Design-time Responsible AI assessments (bias, fairness, transparency); runtime audit logs for all agent decisions, aligned to GDPR, sector-specific regulations
- Shadow agent detection - Technology and process to discover unauthorized AI tools; structured onboarding or retirement procedure for any shadow agent found

Agentic lifecycle management

O

- Design & blueprinting - Every agent starts with a defined purpose, success criteria, and an approved Blueprint; business and IT aligned upfront. No retroactive governance bolt-on.
- Release Readiness & Governance Gate - A governance gate that ensures agents are ready to go live: securely, responsibly, and at scale. It validates approvals, blueprints alignment, identity, security, logging, and policy enforcement, aligned with the client's existing change and release processes.
- Monitoring & performance analytics - Technical KPIs (uptime, error rate) AND business KPIs (queries resolved, time saved, user satisfaction); Monitoring of potential issues. Alters for performance degradation, continuous logs analysis. Tracking licenses usage.
- Feedback & improvement loop - Structured user feedback channels; monthly governance review of agent performance data; mechanism for updates, retraining, or retirement - agents stay "evergreen". Revalidation of templates. Recommendations being share with respective teams.
- Lifecycle Ownership & Accountability - Every agent is governed from activation to retirement; ownership accountability enforced throughout with defined escalation tiers, periodic compliance reviews, and formal decommissioning procedures.

Journey to governed AI - five phases from assessment to scale

Each phase delivers concrete outcomes – from assessment to full enterprise-wide governance program

Assess

Deliverable: Assessment report + maturity heatmap per pillar + high-level roadmap

Conduct Agentic Governance Maturity Assessment: inventory all AI agents including shadow across all hyperscalers, evaluate infrastructure and security, clarify business objectives for AI. Identify which use cases and pain points are driving this initiative.

Plan

Deliverable: Implementation plan / governance charter + technical architecture diagrams

Define governance structures: Agent Owner, AI Steward roles, committees - and enterprise policies for security, compliance, and data usage. Design the target technical architecture integrating Agent 365, Copilot Studio, and Microsoft Foundry, with consistent governance across all hyperscale's and platforms. Align IT, Security, and Business stakeholders.

Deploy

Deliverable: 1–3 governed agents live, control platform running, guardrails active + operational playbooks

Deploy Agent 365 Control Plane; configure Entra ID integration, Microsoft Purview, Defender, Sentinel. Upload initial Agent Blueprints; establish monitoring. Onboard into A365 1–3 pilot agents (high-impact processes) piloted under new governance model; iterate and tune DLP rules and blueprints.

Operate

Deliverable: Operational handbook (incident response, change management) + trained core team

Run governed AI environment for stabilization; fine-tune alerting, log data, support processes. Training workshops for IT admins, support teams, and end-users. governance Establish regular review cadence, monthly AI Ops reviews with KPI tracking.

Scale

Deliverable: Full agentic governance program across agreed scope + future enhancement roadmap + value report

Onboard all agents under governance; apply blueprints across new departments and use cases. Formalize AI Center of Excellence or Governance Board for long-term policy enforcement. Recommend CI/CD enhancements for agent version control. Value report: incidents prevented, time saved, management KPIs achieved.

Start with Phase 1 – Schedule the Maturity Assessment. Low commitment, high impact – the right first step before major investment.

Agentic governance maturity & readiness assessment

Structured diagnostic - the crucial first step before any major investment, ensuring implementation focuses on what matters most

Step 1

Current state review

- Interviews and workshops with IT, Security, and Business stakeholders
- Inventory all AI agents in use (including shadow AI) across all departments and hyperscalers
- Assess infrastructure readiness: Microsoft Foundry, Copilot Studio, M365 licensing, data landscape, existing governance structures
- Gauge organisational readiness: AI strategy clarity, executive buy-in, existing policies and AI governance maturity

Step 2

Maturity evaluation & Gap analysis

- Score maturity using Capgemini's Agentic Governance Maturity Model: from no governance to fully optimized
- Output: heatmap showing strengths vs. weaknesses per pillar - e.g., Platform: do you have a central agent inventory? Security: are DLP controls applied to AI workflows?
- Identify specific gaps per pillar (e.g., shadow AI proliferation = Platform gap; DLP not applied to AI = Security gap)
- Identify strengths to build on (e.g., existing DevOps pipeline that can be leveraged for agent deployment)

Step 3

Roadmap & recommendations

- Phased roadmap aligned to business value and feasibility - prioritized by top risk/opportunity
- Quick wins, such as applying DLP controls to your top AI workflows paired with longer-term milestones that scale governance
- Clear recommendations with each initiative mapped to a maturity phase, a responsible owner, and a measurable outcome
- Executive briefing: maturity heatmap + approved path forward - green light to proceed or immediate actions to take internally

Outcome:

Outcome – Maturity heatmap + gap analysis + phased roadmap delivered to executive leadership - clear, approved plan before major investment. Hitting the ground running in the right direction.

Secure your AI future

Next Steps with Capgemini

Recommended immediate next step: Agentic governance maturity & readiness assessment - a 360° AI governance health check. Know exactly what needs to be done and in what order. Low-commitment, high -impact. Can be scheduled as early as next month with Capgemini experts on-site within weeks.

Next step - begin your governance journey

Looking for improving your operations? Agentic governance maturity & readiness assessment is the first step.

Want to hear more? Schedule a technical Deep Dive sessions.

Are you not sure what agents can do? Let's explore agentic exosystem during a 2-day Hackathon

420,000 + team members globally

50+ years of technology transformation

Official Microsoft agent 365 launch partner

Strategy, assessment, implementation, and ongoing operations - end-to-end partnership

"Make it Real" - Capgemini's commitment: turning AI governance from strategy into tangible results

Beyond governance: Capgemini's unified agentic studio

Industries and activity domains

We can operate in all industries to optimize business processes…

… also corporate activities

Corporate domains

…and implement a robust data and tech foundation to support agentic deployment

Capgemini | Microsoft

Manufacturing	Financial Services	CP & Retail	Life Sciences	Govt/Public
R&d Engineering	Customer onboarding	R&D & product innovation	Research & development	Citizen engagement & services
Manufacturing	Lending & credit	Supply chain	Manufacturing & supply chain	Inspections, permits & compliance monitoring
Supply chain	Payments & transactions	Multi-channel operation	Customer experience	Urban planning & infrastructure
Customer experience	Customer engagement	Manufacturing	Corporate & compliance	Education and workforce development
	Policy management	Customer experience		
	Risk & compliance			
	Claims processing			

Human resources | Finance | Sale | Legal, risks & compliance

Information technology | Procurement | Corporate social responsibility

Ai-ready data & agentic foundation Agentic, AI, data platform, data products, governance all trustable and cybersecure including sovereignty capabilities and with sustainability by design

Agentic workforce in agentic workplace

Capgemini’s guiding hand: Expertise, framework, and tools

Four differentiators that make Capgemini the right partner for your AI governance journey

Official Microsoft launch partner for agent 365

Capgemini is proud to be an official Microsoft launch partner for Agent 365. This strategic partnership ensures early access to Microsoft’s latest innovations. Our teams develop reference architectures, governance frameworks, and deployment accelerators aligned with Microsoft’s roadmap enabling clients to scale AI agents securely and confidently from day one.

01

Raise platform

Capgemini addresses governance end-to-end: Platform (technology) + Security (policy) + Operations (process & people). Full Microsoft ecosystem leverage: Purview, Defender, Sentinel, Entra ID - all aligned with Agent 365. Tools configured to work together, technically sound, secure by design, operationally feasible.

02

Proven playbook

Track record from dozens of enterprise AI transformation projects - industry-specific nuances (banking, retail, healthcare) baked into methodology. Case studies, reference architectures, and deep knowledge of Microsoft’s latest AI capabilities applied to your context.

03

Change empowered with strategy-driven enablement

We don’t just implement technology, we embed change. Capgemini’s approach integrates strategic advisory with hands-on enablement to ensure sustainable adoption of agentic AI. From executive alignment and operating model design to role-based training and CoE setup, we help clients build the capabilities, culture, and confidence to govern and scale AI agents enterprise-wide.

04

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real.
www.capgemini.com

Capgemini