

The hidden risk in your clients' *AI integration strategy*

Model Context Protocol (MCP) is quietly becoming the connective tissue of enterprise AI. Most organisations have no governance framework, no audit trail, and no visibility into what is already connected — or who might be listening.

1,800+

MCP servers on the public internet with no authentication

437k+

downloads of a single package carrying a critical remote code execution flaw

75%

of API gateway vendors projected to support MCP by 2026 — Gartner

THE BASICS

A new universal connector for enterprise AI

You have heard about AI agents. MCP is the layer underneath them that almost nobody is talking about yet.

Launched by Anthropic in November 2024, the **Model Context Protocol (MCP)** is an open standard that allows AI applications to connect directly to external tools and data sources —

email inboxes, files, databases, CRM systems, APIs — through a single, consistent interface. It has been described as *"the USB-C port for AI applications."*

Before MCP, every AI tool required bespoke integrations. Now, a developer can connect an AI agent to an organisation's entire data estate in an afternoon. That is the power of the protocol. It is also the problem.

The analogy that resonates with clients: When REST APIs became ubiquitous in the early 2010s, organisations scrambled to build governance. MuleSoft, Apigee, and Kong became billion-dollar businesses. MCP is the same inflection point — and the window to establish expertise is open right now.

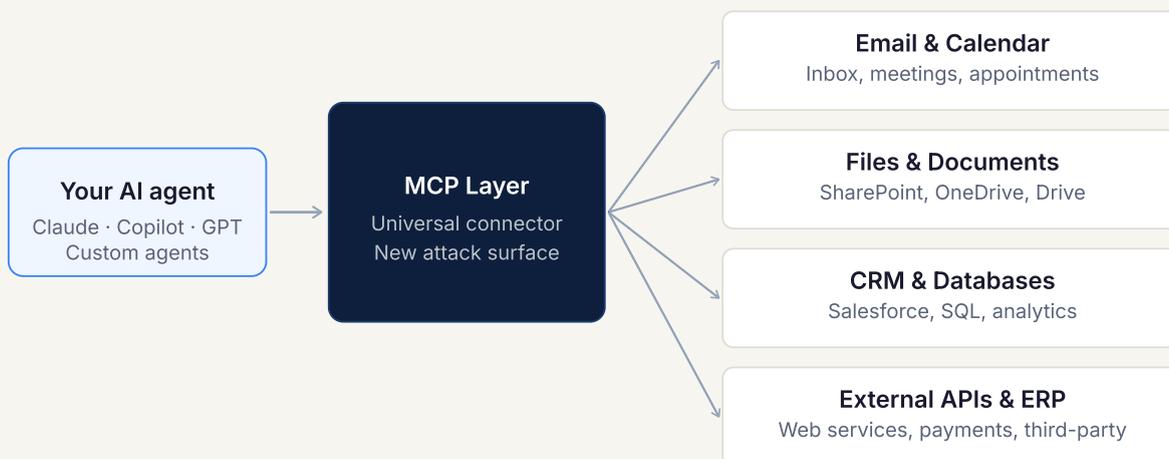


FIG 1 — MCP sits between every AI agent and every enterprise data source. One layer. Every credential.

This is not a future problem. It is a current one.

Your clients are already deploying MCP-connected tools. The majority do not know it is happening.

75%

of API gateway vendors will include MCP features by 2026 — making it structurally unavoidable across most enterprise software stacks.

Gartner, Software Engineering Hype Cycle, 2025

33%

of enterprise software projected to include agentic AI by 2028, up from less than 1% today — meaning the MCP attack surface grows exponentially.

Gartner, 2025 Software Engineering Survey

<18mo

since MCP was publicly released (November 2024). Authentication was not even part of the original specification — OAuth was only added in March 2025.

Anthropic MCP Specification changelog, 2025

"MCP's rapid adoption and continuous evolution are creating security risks that product leaders focusing on securing AI must address."

— Gartner, *Emerging Tech: Security Implications of Model Context Protocol*, 2025

When Microsoft integrated MCP support across Copilot Studio, Azure AI Foundry, and the broader Microsoft 365 ecosystem, the enterprise deployment timeline collapsed. Your clients' developers are already adding MCP servers as productivity tools — not as IT infrastructure decisions. There is almost certainly no approval process, no register of what is connected, and no monitoring.

THE THREAT LANDSCAPE

Five ways MCP becomes a liability

MCP's power — direct, programmatic access to enterprise data — is identical to its risk. Once an AI agent can read and write to your client's systems, so can anyone who can influence that agent.



Tool poisoning

A malicious MCP server embeds hidden instructions inside its own tool description. When the AI reads it, the instructions execute silently — triggering file reads, email forwards, or data transfers the user never authorised.

OWASP LLM Top 10 – #1 Risk



Indirect prompt injection

An attacker embeds a hidden command inside a document, email, or web page. When an AI agent reads that content via MCP, it follows the attacker's instructions — not the user's. Zero user interaction required.

CVE-2025-32711 – Microsoft 365 Copilot



Supply chain compromise

Popular MCP packages can be modified after installation. In 2025, a widely-used email MCP server was updated to silently BCC all outbound emails to an attacker's address. Users saw nothing unusual.



Keys to the kingdom

MCP servers typically hold authentication tokens for every connected service. Compromise one server and an attacker gains persistent access to email, files, CRM,

CVE-2025-6514 – 437,000+ downloads

and databases — without triggering a password-change alert.

High-value target for lateral movement



No authentication · no audit trail

Authentication was absent from the original MCP specification. Over 1,800 servers are live with no auth required. Most implementations lack standardised logging — making forensic investigation nearly impossible.

Knostic research, 2025

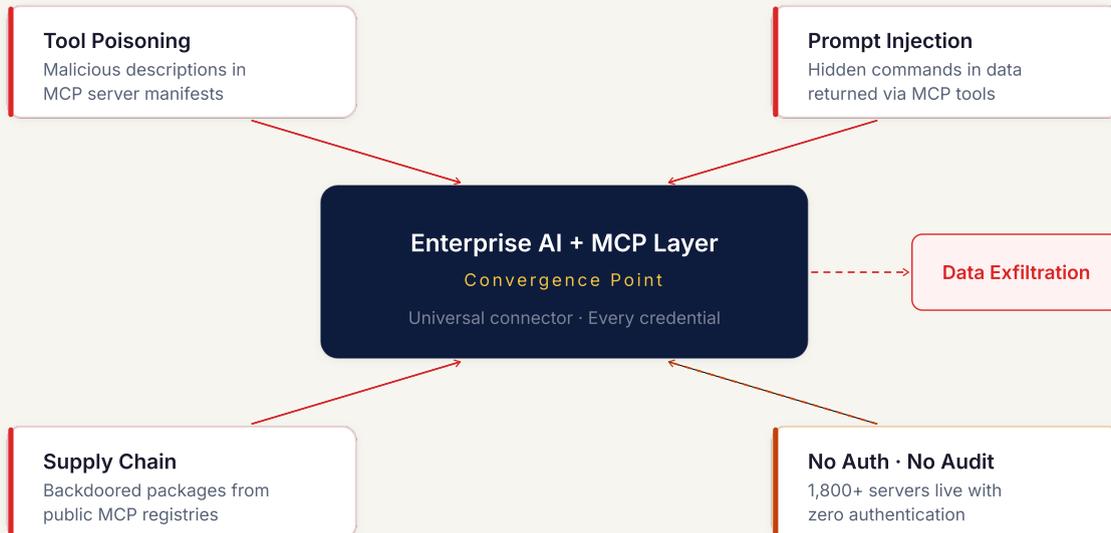


FIG 2 — MCP's position as the universal connector makes it a high-value convergence point for multiple attack vectors simultaneously

Real incidents. Real organisations. Real exposure.

MCP is less than 18 months old. There is already a documented timeline of security failures. These are not theoretical.

2025 — POSTMARK MCP SERVER

Silent email exfiltration via routine package update

A widely-used MCP server for email sending — with over 1,500 weekly downloads — was modified to add a hidden BCC field. Every outbound email sent by every AI agent using this server was silently copied to the attacker's address. Users saw no indication anything had changed. The attack vector: a malicious maintainer pushing a routine-looking update.

2025 — MICROSOFT 365 COPILOT **CVE-2025-32711**

EchoLeak: zero-interaction data exfiltration via hidden prompts

Threat actors discovered that embedding tailored hidden prompts within a Word document or email was sufficient to cause Copilot to silently exfiltrate sensitive data. The victim needed only to ask Copilot to summarise the file — no clicks, no confirmations, no alerts. The attack exploited MCP's indirect prompt injection vulnerability at enterprise scale.

2025 — MCP-REMOTE OAUTH PROXY **CVE-2025-6514 · CVSS 9.6**

Critical remote code execution in a package with 437,000+ downloads

A critical OS command-injection vulnerability in mcp-remote — the standard OAuth proxy for connecting MCP clients to remote servers — allowed a malicious MCP endpoint to execute arbitrary commands on the victim's machine. The package was referenced in integration guides from Cloudflare, Hugging Face, and Auth0. Any unpatched install was a supply-chain backdoor: API keys, cloud credentials, SSH keys, and repository contents could all be extracted.

2025 — ASANA MCP

Cross-tenant data exposure across enterprise customer instances

A bug in Asana's MCP implementation caused data to leak between different enterprise customer instances. An MCP bug in one of the world's largest project management platforms resulted in data from one organisation being accessible to another — a breach that would trigger regulatory notification obligations for most regulated entities.

2025 — WHATSAPP MCP SERVER (PROOF OF CONCEPT)

Entire message history exfiltrated via tool poisoning

Security researchers demonstrated that by combining a legitimate WhatsApp MCP server with a malicious tool description in the same agent, they could silently exfiltrate a user's complete WhatsApp message history. The attack required no user action beyond asking the AI a routine question. The vulnerability is inherent to any MCP deployment without tool verification and allowlisting.

"The MCP registry requires only proof of GitHub repository or domain ownership — it does not require code review, security audit, or malware scanning. A server listed in the official registry is no more trustworthy than any other community package, yet users may incorrectly assume registry presence implies vetting."

— arXiv: Securing the Model Context Protocol (MCP), 2025

THE GOVERNANCE GAP

Why your clients almost certainly do not know this problem exists

MCP is arriving in organisations through the front door — developers installing it as a productivity feature in their IDE, their AI assistant, their copilot — not through IT procurement. This is the **shadow AI** problem, and it is structurally different from shadow IT of the past.

Shadow IT was visible on the network. Shadow MCP is embedded inside tools that already have IT approval. A developer adding an MCP server to Claude Code or GitHub Copilot does not trigger a firewall alert, a procurement review, or an access control workflow. It simply works.

The result: organisations are building MCP infrastructure right now without governance, without an inventory of what is connected, and without the ability to answer the question a regulator or a CISO will eventually ask — *"which AI agents have access to which data, and how do you know?"*

The five reasons clients are unprepared

1. MCP is less than 18 months old — it predates existing AI security policies in virtually every organisation.
2. It arrives via developer tooling (IDEs, copilots), bypassing IT procurement and change control entirely.

3. There is no standard MCP server register, approval workflow, or audit log format — nothing to flag a review.
4. The official MCP registry provides a false sense of security — listing requires no security audit whatsoever.
5. Security teams are focused on LLM data inputs — not on the tool execution layer underneath.

The parallel that resonates

In 2012, REST APIs were everywhere and almost nobody had an API security programme. Companies were exposing internal data through public-facing endpoints without rate limiting, authentication audits, or centralised visibility.

The API governance market (MuleSoft, Apigee, Kong, Tyk) emerged to solve exactly that problem and grew to tens of billions in value.

MCP is the 2025 equivalent. The protocol is already deployed. The breaches are already happening. The governance market is not yet built.

For AheadMG: The firms that build MCP governance capability now will own a category that is about to become mandatory for every enterprise running AI agents. The clients who move first will have competitive advantage. Those who move last will have regulatory exposure.

FCA-regulated clients face compounding risk. DORA and the FCA's operational resilience rules require firms to map, test, and govern critical third-party dependencies. An ungoverned MCP layer — connecting AI agents to customer data, trading systems, or financial records — is a compliance exposure as well as a security one. Regulated clients have both the motivation and the obligation to address this.

A three-stage managed service for MCP governance

The opportunity is not to build gateway software — open-source tools and commercial platforms already exist. The opportunity is to own the governance and managed service layer on top, packaged as a repeatable offering AheadMG can sell across its client base.

STAGE 1 — DISCOVERY

1

MCP Risk Assessment

Inventory all MCP servers deployed or planned across the client estate. Assess each against key threat vectors. Identify governance gaps. Deliver a board-ready risk briefing and a prioritised remediation roadmap.

Deliverable: Risk briefing + MCP server register + CISO presentation

4 – 6 weeks

STAGE 2 — FRAMEWORK

2

Governance Framework Deployment

Implement policy-as-code for MCP server approval workflows. Deploy RBAC and least-privilege access. Integrate with the client IdP (Entra ID / Okta). Configure audit logging. Build the change control process for new MCP server registration.

Deliverable: Deployed governance framework + runbooks + training

8 – 12 weeks

STAGE 3 — MANAGED SERVICE

3

Ongoing MCP Governance Retainer

Monthly MCP server review board (approve / reject new servers). Continuous monitoring dashboards and anomaly alerting. Quarterly CISO-ready governance report. Incident response support for MCP-related events.

Deliverable: Retained managed service with named CISO accountability

The FEAW Services angle: FEAW can provide the technical deployment capability — MCP gateway configuration (Bifrost/MintMCP), devcontainer and CI/CD integration, Entra ID wiring, audit log pipelines — while AheadMG owns the client relationship, the governance framework design, and the managed service commercial model. A clean division of labour that creates a differentiated joint offering neither firm could deliver alone.

THE CLIENT CONVERSATION

Six questions to open the conversation with any client

Most clients will not know this problem exists. These questions are designed to surface the gap without requiring them to know what MCP is first.

1 "Do you have an AI agent programme, or AI tools connected to your business systems?"

Establishes whether MCP is likely in scope. If they say yes — to Microsoft Copilot, GitHub Copilot, Claude, ChatGPT Enterprise, or any custom agent — MCP is almost certainly already deployed.

2 "Who in your organisation is responsible for approving connections between AI tools and your data sources?"

Identifies the governance gap immediately. The answer is almost always silence, or a name that is followed by "but I'm not sure they know about all of them."

3 "Do you know which of your cloud services — Microsoft 365, Salesforce, ServiceNow — are currently connected to AI agents?"

Reveals the shadow inventory problem. Creates the moment of realisation that there may be connections nobody has mapped.

4 "What is your current process for approving a new integration between an AI tool and business data?"

Uncovers the absence of process. If there is no answer, the governance gap is confirmed. If there is an answer, the follow-up is: "And does that cover MCP servers installed by individual developers?"

5 "Have you reviewed your AI tool supply chain since the Microsoft 365 Copilot vulnerabilities were disclosed last year?"

Introduces urgency using a known, credible incident (CVE-2025-32711). References a vendor they definitely use. Makes the risk concrete rather than theoretical.

6 "Is MCP governance covered in your current AI security or acceptable use policy?"

The answer is almost certainly no — the protocol is too new and the policy landscape has not caught up. This creates a clear entry point: "We can help you build that framework."

Recommended next steps for AheadMG

Survey two or three existing clients using questions 1–3 above. Gauge whether the governance gap exists and whether there is appetite for a structured response. Use the findings to validate the service design before building the full offering.

[Download client survey template](#)

[Read the Verity platform proposal](#)

Sources and references

Gartner (2025). *Emerging Tech: Security Implications of Model Context Protocol*. [gartner.com](#) — Gartner (2025). *Software Engineering Hype Cycle: MCP Adoption Projections*. [gartner.com](#) — K2View (2025). *MCP Gartner Insights*. [k2view.com](#) — arXiv (2025). *Enterprise-Grade Security for the Model Context Protocol (MCP): Frameworks and Mitigation Strategies*. [arxiv.org/abs/2504.08623](#) — arXiv (2025). *Securing the Model Context Protocol (MCP): Risks, Controls, and Governance*. [arxiv.org/abs/2511.20920](#) — eSentire (2025). *Model Context Protocol Security: Critical Vulnerabilities Every CISO Must Address*. [esentire.com](#) — AuthZed (2025). *A Timeline of Model Context Protocol Security Breaches*. [authzed.com](#) — Red Hat (2025). *Model Context Protocol (MCP): Understanding Security Risks and Controls*. [redhat.com](#) — Pillar Security (2025). *The Security Risks of Model Context Protocol*. [pillarsecurity.com](#) — Bitdefender Business Insights (2025). *Security Risks of Agentic AI: A Model Context Protocol Introduction*. [bitdefender.com](#) — Knostic (2025). *MCP server authentication research*. [knostic.ai](#) — CVE-2025-32711: Microsoft 365 Copilot EchoLeak vulnerability. [NVD](#) — CVE-2025-6514 (CVSS 9.6): mcp-remote OS command injection. [NVD](#).

Prepared by FEAU Services Ltd for AheadMG. March 2026. Confidential — not for distribution.