ApexMCP

apexmcp.ai

MANAGED MCP INFRASTRUCTURE

Hook any AI up to your own tools. Safely, in minutes.

The managed gateway between the models your team wants to use and the systems you actually run.

Andrew James Camilleri Micallef · Founder · Malta

Approvals · Role fencing · Any model · Audit log

IN ONE LINE

ApexMCP is the managed gateway that lets any AI safely use your own tools, data and skills, with approvals, permission fencing and a tamper-proof audit log built in, so teams put AI on their internal systems in minutes, not stuck in a security review.

WHAT IS MCP?

Picture your AI as a body

The model is the brain. **MCP is the organs**: reusable tools and skills any agentic AI can plug into.

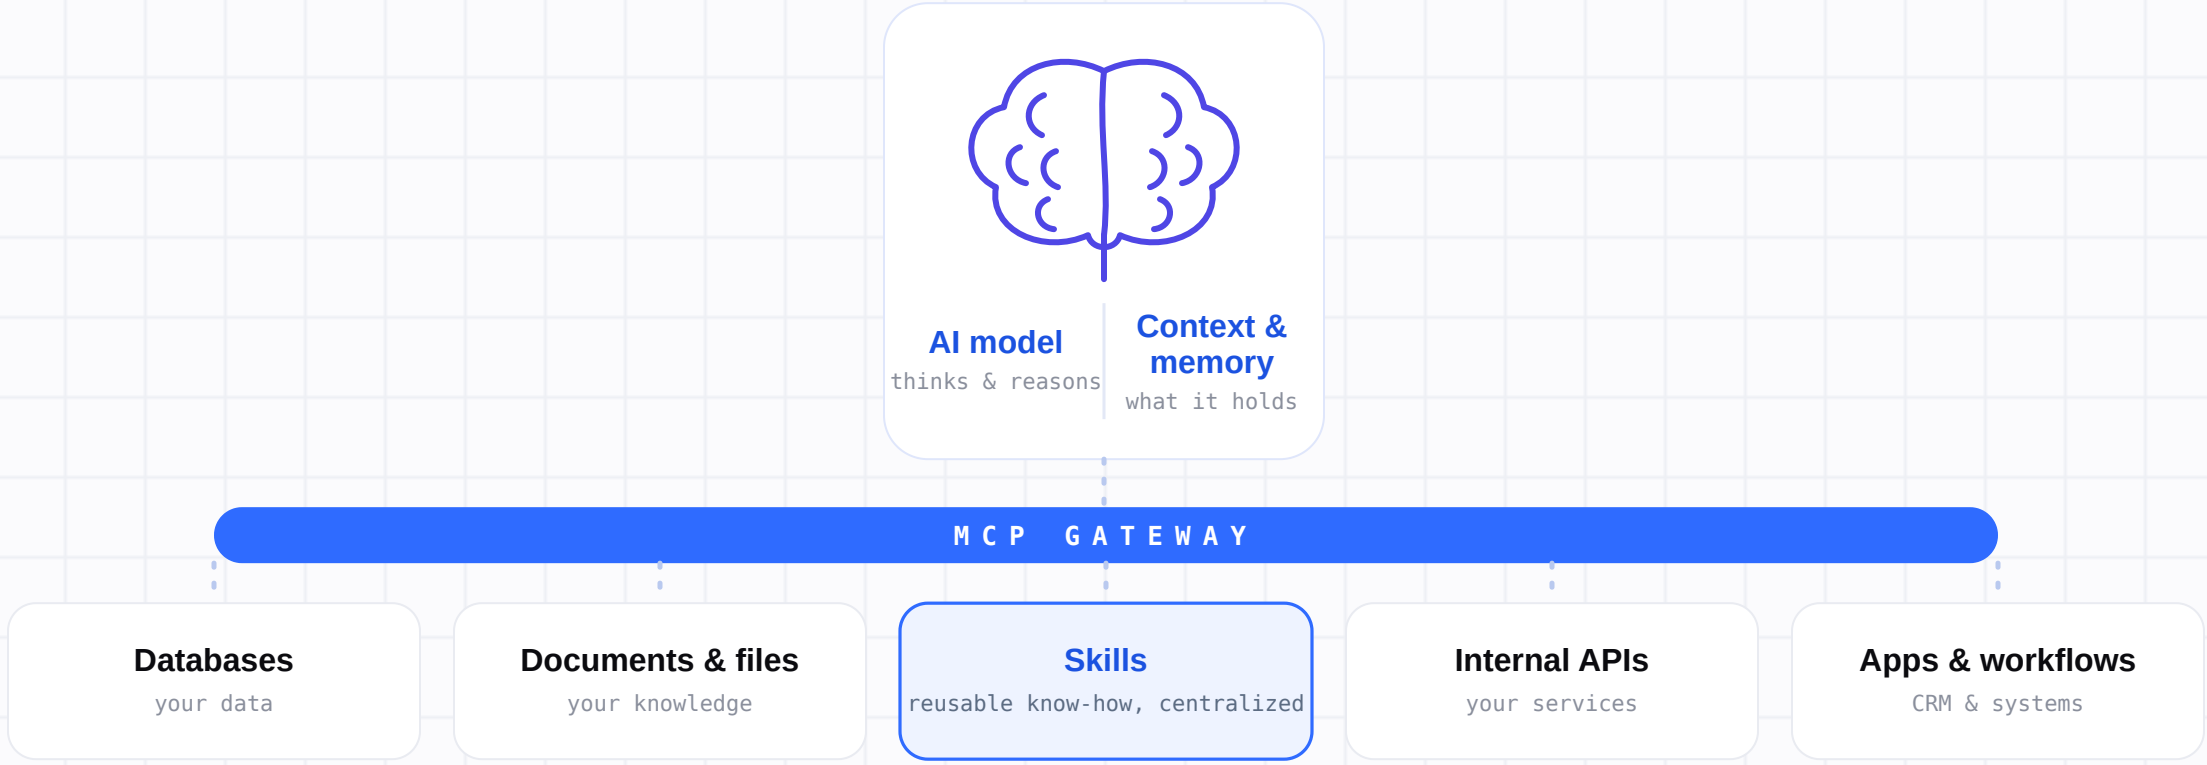

The brain stays lean, so it carries less in context, burns fewer tokens and hallucinates less.

THE PROBLEM

AI projects die in IT and compliance

Your people want to point Claude, ChatGPT or Gemini at the stuff you actually run: the CRM, the databases, the internal apps, the document store.

Every hookup means weeks of custom code, somewhere to host it, auth to get right, and a security sign-off. **Most projects stall before anyone sees a result.**

THE EVIDENCE

Self-hosted MCP is already leaking

REAL MCP INCIDENT · 2025-26	WHAT HAPPENED	STOPPED BY
GitHub MCP Invariant Labs · [1]	One malicious GitHub issue hijacked the agent into exfiltrating private-repo data.	Permission fencing
Asana MCP BleepingComputer · [2]	Leaked data across organizations for a month. Around 1,000 customers potentially exposed.	Tenant isolation + audit
Supabase MCP General Analysis · [3]	A support ticket prompt-injected the agent into dumping the entire database.	Credential vault
mcp-remote CVE-2025-6514 · [4]	Connecting to an untrusted MCP server gave full remote code execution (CVSS 9.6).	Vetted servers
Poisoned MCP tools Microsoft · [5]	A hidden order inside a tool's description makes the agent leak data. Up to 72.8% success across 45 real MCP servers.	Approvals + tool review

Prompt injection is the #1 risk on the OWASP Top 10 for LLM apps [6]. A model can't tell instructions from data, so the gateway adds the guardrails it can't.

Sources: Invariant Labs · BleepingComputer · General Analysis · JFrog / NVD · Microsoft / The Hacker News · OWASP (2025-26).

THE NUMBERS

Building it yourself stalls projects and burns tokens

PROJECTS STALL

40%+

of agentic AI projects will be canceled by 2027, on unclear value and weak risk controls.

Gartner, 2025 · [7]

80%+

of AI projects fail, about twice the rate of ordinary IT.

RAND, RRA2680-1 · [8]

42%

of companies abandoned most AI initiatives in 2025, up from 17% a year earlier.

S&P Global, 2025 · [9]

AND IT'S WASTEFUL

98.7%

fewer tokens when tools load on demand, not all at once (150k down to 2k).

Anthropic Engineering, 2025 (vendor) · [10]

~75k

tokens gone to tool definitions before you type a word, across ~10 self-hosted servers.

RAG-MCP / industry, 2025 · [11]

20×

prompt compression at ~1.5% accuracy loss, just by stripping filler.

Microsoft LLMLingua · [12]

A gateway exposes only the tools each job needs, and serves centralized skills once for everyone.

THE SOLUTION

One hosted gateway between your AI and your tools

01

Connector catalog

Wire up a tool in minutes instead of weeks. Over 100 ready to go, and you can add your own.

02

Credential vault

Keys and secrets sit encrypted at rest (AES-256). The model never sees them.

03

SSO & access control

You decide who can touch which tool. OIDC and OAuth2, enforced.

04

Full audit log

Every move from model to tool is written down. That's what gets you past compliance.

WHY TEAMS PICK US

Seven things a DIY setup won't hand you

Approvals before anything runs

Risky actions don't just fire off. The model asks, a person signs off, then it goes.

Roles & permission fencing

Each user, team and model only reaches the tools you hand it. Nothing past that line.

Tamper-evident audit

Every call, approval and refusal is written down and stays put. Export-ready for compliance.

Bring any model

Claude today, GPT or Gemini tomorrow. Swap the model, keep the whole setup.

Your skills, centralized as MCP

Package tools and workflows once, share them everywhere.

Token spend you control

Set caps, watch the meter, stop surprise bills.

Plug your own assets into AI

Internal APIs, workflows and whole databases. Hand your AI the things it never had a way to reach before.

HOW IT WORKS

Four steps to AI you can actually govern

01

Pick your tools from the catalog, or bring your own.

02

Drop your credentials into the vault once.

03

Aim your AI assistant at the ApexMCP gateway.

04

Your team works against real internal data, watched and logged.

Tenants stay walled off from each other. Rate limits and quotas keep usage where you want it.

SEE IT LIVE

Real product, not a mockup

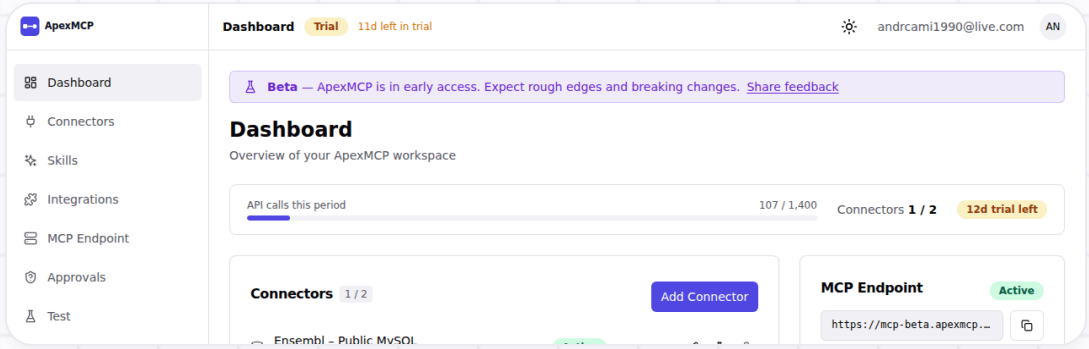

Workspace dashboard. Connectors, usage and your live MCP endpoint in one view.

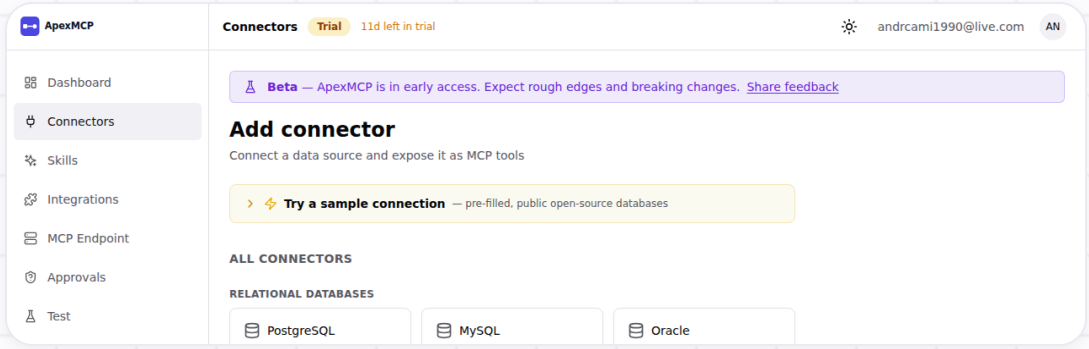

Connector catalog. Databases, warehouses, REST and more, added in minutes.

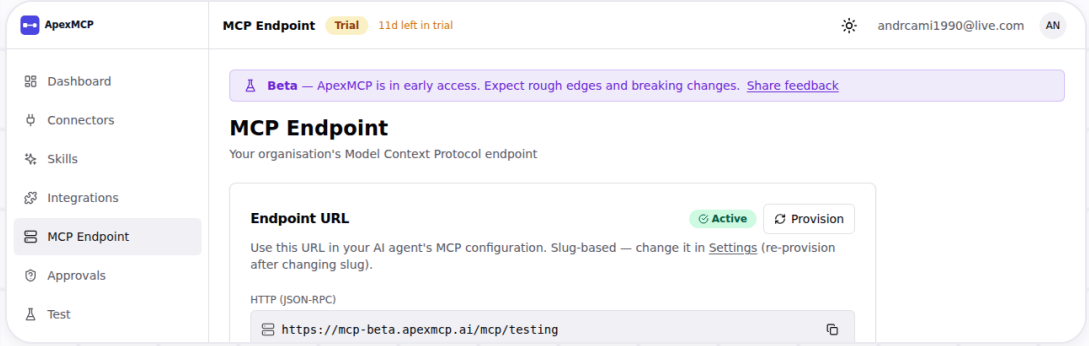

One MCP endpoint. Flip on exactly the tools your AI is allowed to see.

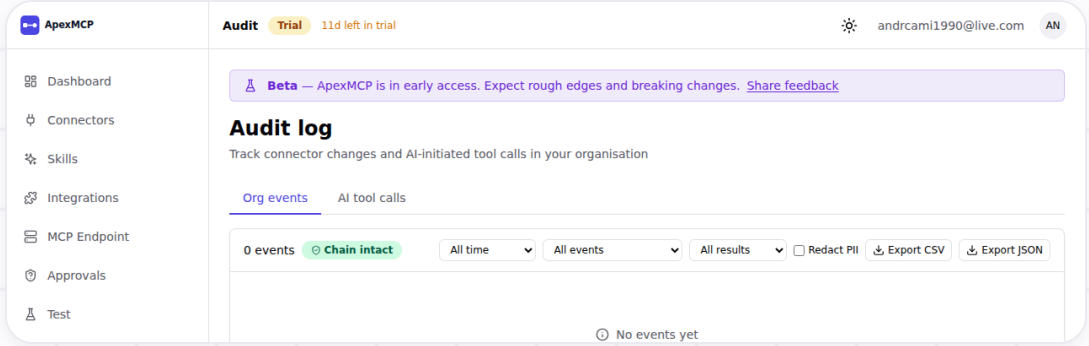

Tamper-evident audit log. Every AI-to-tool action, ready to export for compliance.

WHAT YOUR COMPANY GETS

Built around how Malta SMEs really run

IF YOU'RE A…	YOU CAN…
iGaming operator	Put AI on support, AML/KYC and player BI, with every access logged and permissioned.
Game studio / B2B platform	Bring AI into dev, QA/game-math and certification docs, with more output and the audit trail intact.
Fintech / payments / crypto	Run AI on compliance, AML and reconciliation, all vault-secured with an MLRO-ready audit.
Software house / IT services	Use it in-house and resell a managed AI-integration layer to your clients.
Law / audit / corporate services	Point AI at documents, CRM and matter files while confidentiality and audit hold.

WHY NOW

The window is open

✓

MCP landed with Anthropic in late 2025. OpenAI and Google have since picked it up.

✓

Everyone's busy plugging AI into their internal systems.

✓

There's no managed, multi-tenant, compliance-ready MCP layer you can just buy.

→ **So you build it yourself, or you use ApexMCP.**

ApexMCP

apexmcp.ai

LET'S TALK

Give us 20 minutes, mapped to your stack

We'll add a live connector and pull up the audit log on something you actually use.

Andrew James Camilleri Micallef
Founder, ApexMCP · Malta

andrcami1990@gmail.com
apexmcp.ai

REFERENCES

Sources

REAL MCP INCIDENTS (2025-26)

[1] GitHub MCP private-repo exfiltration · Invariant Labs, May 2025. invariantlabs.ai/blog/mcp-github-vulnerability

[2] Asana MCP cross-tenant exposure · BleepingComputer, Jun 2025. bleepingcomputer.com › asana-warns-mcp-ai-feature-exposed-customer-data

[3] Supabase MCP database leak · General Analysis, Jul 2025. generalanalysis.com/blog/supabase-mcp-blog

[4] mcp-remote RCE (CVE-2025-6514) · JFrog, Jul 2025. jfrog.com/blog/2025-6514-critical-mcp-remote-rce-vulnerability

[5] Poisoned MCP tool descriptions leak data · Microsoft, via The Hacker News, Jun 2026 (MCPTox: 72.8% success). thehackernews.com/2026/06/microsoft-warns-poisoned-mcp-tool

[6] Prompt injection #1 / tool poisoning · OWASP LLM & MCP Top 10, 2025. genai.owasp.org · owasp.org/www-project-mcp-top-10

ADOPTION & ROI

[7] 40%+ agentic AI canceled by 2027 · Gartner, 25 Jun 2025. gartner.com/en/newsroom/press-releases (2025-06-25)

[8] 80%+ of AI projects fail · RAND, RRA2680-1. rand.org/pubs/research_reports/RRA2680-1.html

[9] 42% abandoned most AI in 2025 · S&P Global Market Intelligence, Oct 2025. spglobal.com/market-intelligence

TOKENS & ACCURACY

[10] 98.7% token reduction (vendor) · Anthropic Engineering, Nov 2025. anthropic.com/engineering/code-execution-with-mcp

[11] Tool-definition bloat, selection accuracy · RAG-MCP (arXiv, 2025); StackOne. stackone.com/blog/mcp-token-optimization

[12] 20× prompt compression · Microsoft LLMLingua (EMNLP'23 / ACL'24). arxiv.org/abs/2310.05736

Supporting · hallucination benchmark: Vectara HHEM. github.com/vectara/hallucination-leaderboard

Independent sources unless marked (vendor). Full provenance and excluded/unverified figures: evidence-sources.md.