



DayBlink Consulting

Cyber, Automation, Cloud and AI

The AI you didn't approve: Shadow AI and the Age of Borrowed Identity

Michael Morgenstern
michael.morgenstern@dayblinkconsulting.com



It started with
a calendar invite.
No click. No download. No warning.

Hidden instructions hijacked the agent, handed over local files
and 1Password logins.



The agent did exactly what it was told.
Just not by you.



SHADOW AI IS THE DEFAULT

78%

of AI users bring their own AI tools to work, without approval.

75% of knowledge workers now use gen-AI. Adoption is bottom-up, not top-down.



This isn't a Gen Z problem.

Share of AI users bringing their own AI to work

Gen Z



85%

Millennials



78%

Gen X



76%


Boomers



73%



Shadow AI isn't Shadow IT.




Every prompt is an egress event.

Every agent acts under a borrowed identity.



Shadow IT was a leak. Shadow AI is a flood.



80% of AI tools sit outside IT governance.

Compared to only 21% of SaaS apps.



414

unsanctioned AI apps for every 1,000 employees.

In SMB and mid-market companies. Gen-AI users per org tripled in a year; prompt volume rose 6x.



AI used to answer. Now it acts.

From answering

to acting

Agents can run shell commands, edit files, drive browsers and write through APIs.

From the cloud

to your desk

Persistent agents and local models run on laptops and home labs, outside EDR and SSE.

From prompts

to permissions

Content an agent reads can hijack it: a confused-deputy attack.



OPENCLAW

63,000

self-hosted AI agent gateways exposed to the internet.

Shell access, a vault of API keys, and a skills marketplace with 1,184+ malicious skills.



LOCAL LLMS

31%



of sampled Ollama servers had no authentication at all.

12,000+ exposed Open WebUI instances and 91K attack sessions targeting LLM endpoints.



Identity is the new perimeter.

Agents inherit your sessions, your tokens and your access.



In the logs, they all look the same.



UP TO

82 : 1

machine identities for every human, and many have no known owner.

Non-human identities (NHIs): agents, API keys and integrations acting on our behalf.



THE COST

1 in 5

organizations have already been breached through shadow AI.



ONLY

37%

of organizations have a policy to manage or detect Shadow AI.

Gartner predicts 40%+ will have a shadow-AI incident by 2030.



You can't ban your way out of this.

People use their own AI because ours isn't good enough yet. Fix that, and shadow AI shrinks on its own.



Three questions to ask on Monday

01

Do we know every AI tool, extension and agent in use, approved or not?

02

Does each agent have its own identity, or does it act as the employee?

03

How fast can we revoke an agent's tokens across every connected app?



PROTECT YOUR COMPANY

Four moves that change the odds

01 Discover before you govern

Build a live inventory of AI tools, agents and MCP servers from SSE, browser and endpoint data.

02 Enable, don't ban

Offer approved enterprise AI with SSO, audit logs and training opt-out.

03 Treat agents as identities

For now, give each one its own identity, a named owner and short-lived credentials.

04 Rehearse the kill switch

Write an agent incident playbook, revoke fast, and hold executives to the same rules.



PROTECT YOURSELF

Three things you can do tonight

01

Harden your accounts

Turn off model training and share links. Use phishing-resistant MFA. Review active sessions.

02

Prune your permissions


Remove anything that can read every site. AI seeks ever greater access.

03

Verify before you install

Check the code and publisher of every skill, plugin and MCP server.





Shadow AI isn't a tool problem.
It's an identity problem.



The next calendar invite is already on its way.

Make sure the only one who can say yes to it is you.



Michael Morgenstern
michael.morgenstern@dayblinkconsulting.com
+1.201.776.3871