clutch

A Sobering Look at AI Agents in the Enterprise

EXECUTIVE SUMMARY

The Agent Workforce Is Already the Majority

AI agents outnumber the engineers they assist, hold real credentials, and grow 29% a month — governed by no one.

Clutch Security analyzed observed AI coding agent activity across **tens of live enterprise environments** — spanning **millions of endpoints** and every major sector — over the 90 days ending July 2026. Four findings stand out.

15

Agents are credential aggregators

The average agent uses 15 distinct credentials — keys, tokens, and passwords — more than the person in front of it ever will. Compromise the agent and you inherit everything it holds.

Do This → Organizations should maintain a live inventory of which credentials each agent can reach, enforce least privilege on that access, and continuously monitor how each credential is used.

78%

Agents already operate in production — mostly with the brakes off

Agents touched production systems in 78% of enterprises, and 83% of them had agents running with every approval prompt disabled.

Do This → Organizations should treat agent access to production as privileged access — define a policy for what agents may touch, and enforce it in agent runtime.

73%

MCP governance covers the minority of agent activity

73% of agent tool use bypasses MCP entirely — agents drive credentialed CLI tools directly.

Do This → Organizations should govern agents at the endpoint and credential layer; an MCP gateway alone misses three-quarters of what agents actually do.

3.6x

The agent workforce is already the majority — and compounding

3.6 working agents per engineer, growing 29% a month. One ordinary developer laptop held four agents wielding 88 credentials into production databases and cloud IAM.

Do This → Organizations should run the census now — which agents exist, on which machines, holding what — then define which agents are sanctioned and hold the rest to that bar.

01. THE CENSUS

Agent Adoption Is Universal and Unmanaged

The average enterprise runs a dozen distinct AI agent products and approved none of them.

Ask a CISO which AI agents run in their environment and the answer is a list of what was approved. The approved list and the running list are different documents. The average organization we studied runs a **dozen distinct agent products**, and a large enterprise carries agent installations in the thousands — developers install what makes them productive, and the enterprise inherits the sum of those decisions.

16%

of monitored endpoints show agent activity in the median enterprise — up to 69% in the most saturated

6.3%

of agent-bearing endpoints already run three or more distinct agents

6/10

environments have employees building their own assistants — Copilot Studio, SharePoint, Agent Builder

The market, as it actually deployed

GitHub Copilot, Claude Code, and Cursor are effectively universal — each present in more than nine of every ten environments. Behind them sits a long tail. The enterprise did not standardize on an agent; it accumulated them. The bar shows the share of environments where each agent is present.

The Bottom Line

Nine in ten enterprises run Claude Code, Codex, and Copilot side by side — three vendors billed for one job, on the same endpoints, each with its own credentials. The redundancy isn't just spend; it multiplies the surface an attacker inherits from a single laptop.

02. THE TRAJECTORY

The Agent Population
Doubles Every Ten Weeks

Active agents compound 29% a month, measured from behavior — not our rollout. The curve is not slowing.

Active agents grew a median of **29% month over month** across the environments with enough history to measure it — and nearly every one of them is growing. Indexing each environment to its own starting point and averaging removes the effect of size: the shape below is pure growth, and it doubles the working-agent population about every ten weeks.

4.5x

growth in active agents over six months (indexed to 100)

a median of +29% every month · doubling ≈ every ten weeks

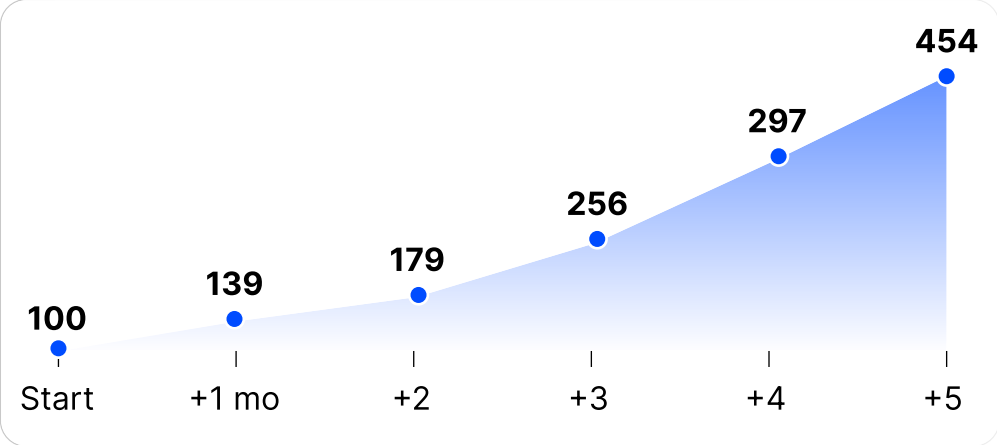

Employees are building agents, not just running them

In **6 of every 10 environments**, employees construct their own assistants — Copilot Studio bots, SharePoint agents, Agent Builder creations — wired to internal data and left running. Governance written for purchased software does not see homemade software.

The autonomous frontier is already inside

Viral autonomous agents — OpenClaw and its family of forks — appear in roughly **2 of every 3 environments**, months after release. Adoption no longer waits for procurement or a security review. It waits for a single install command.

Honorable mentions · also detected in the wild, at negligible volume

Aider · Amp · Cline · Continue · Zed · OpenHands · Kilo Code · Tabnine · Factory · Amazon Q · AnythingLLM · Goose · Antigravity · Lingma · Qod · Manus · Hermes · Warp · Kiro · OpenClaw

The Bottom Line

Whatever you count today is larger by next quarter. Today's negligible agent is next quarter's Claude Code — and the governance gap widens with every doubling.

03. THE RATIO

Agents Outnumber Engineers 3.6 to 1

In 84% of enterprises the machines already hold the majority of the technical workforce.

3.6:1

Working AI agents for every technical employee, in the median enterprise. Not installed — working, with observed activity in the last 90 days. In **84% of environments** agents already outnumber the technical staff they operate alongside, and in the leaders the ratio runs past 37 to 1. Look across the whole workforce and there is roughly one working agent for every four to five employees today — the floor, climbing 29% a month.

Consider what this inverts. Every human in that denominator was interviewed, background-checked, onboarded with provisioned access, and will be offboarded on departure. The agents in the numerator arrived through a single install command — no interview, no joiner-mover-leaver process, no deprovisioning. The entire identity program was built around the smaller population.

Working Agents per Technical Employee · by Sector

Technology	Retail & Consumer	Media & Entertainment	Financial Services	Manufacturing & Industrial
3.8x	3.7x	3.5x	3.4x	0.7x

Agents have saturated engineering at three to four per head across every knowledge-work sector — and heavy industry is the lone laggard, for now. The pattern holds regardless of company size or region.

The Trajectory: Five per Employee by 2027

Hold the observed rate and the arithmetic is severe. Agents reach parity with the entire workforce — one working agent for every employee, not just every engineer — within about six months, by early 2027. Within a year the population grows more than twentyfold: roughly five AI agents for every employee and more than eighty per engineer. Growth this steep must eventually meet saturation — but nothing in the data suggests it is close.

The Bottom Line

The agent workforce is on track to become the largest identity population in the enterprise before the current planning cycle ends — and it is the one population no identity program governs.

04. THE IDENTITIES

Each Agent Holds Fifteen Credentials

A single credential-touching agent aggregates fifteen distinct non-human identities — more keys than the human beside it.

In the environments studied, AI agents used **a median of more than a thousand distinct non-human identities per enterprise** over 90 days. The distribution matters more than the total: **a single credential-touching agent uses a median of 15 distinct non-human identities**, and the busiest do far more — at more than one enterprise a single agent was observed using tens of thousands of distinct credentials, in one case over **67,000**. Each agent is a credential aggregator, one process quietly holding more keys than the person in front of it ever will.

What the identities unlock

29% Source-control tokens	**21%** Cloud infrastructure keys
20% AI-service credentials	**18%** Passwords & static secrets
7% Collaboration & SaaS tokens	**4%** Database credentials

Where Identities Live

GitHub	**98%**
AWS	**98%**
Atlassian	**96%**
PostgreSQL	**93%**
Azure	**91%**
Google Cloud	**91%**
Databricks	**87%**

Agents use live credentials within the hour

A new human employee waits days for access to be provisioned. A new agent does not wait. **One in fourteen agents touched a live credential within its first hour of existence**. Within 24 hours, one in seven had. The agent inherits whatever the endpoint already holds — the AWS profile in the home directory, the GitHub token in the config file, the database password in the environment variable — the moment it starts, with no approval step in between.

7.2% → of new agents used a real credential within 1 hour of first appearing

14.7% → did so within 24 hours — measured from first observed activity

15% distinct identities held by the median credential-touching agent

The Bottom Line

Every agent is a credential aggregator that no IAM program provisioned, monitors, or can revoke. Compromise the agent and you inherit its entire keyring at once.

05. THE TOOLS

Three-Quarters of Agent Activity Bypasses MCP

Agents drive the command line directly. The protocol everyone is watching sees the minority of traffic.

MCP — the Model Context Protocol — has become the headline story of agent tooling. The data tells a quieter, more surprising one. Agents reach infrastructure mostly the old-fashioned way: they drive the credentialed command-line tools already on the endpoint. git, gh, aws, psql, kubectl, docker, ssh — configured long before the agent arrived, and inherited the moment it starts.

How agents invoke tools · median enterprise

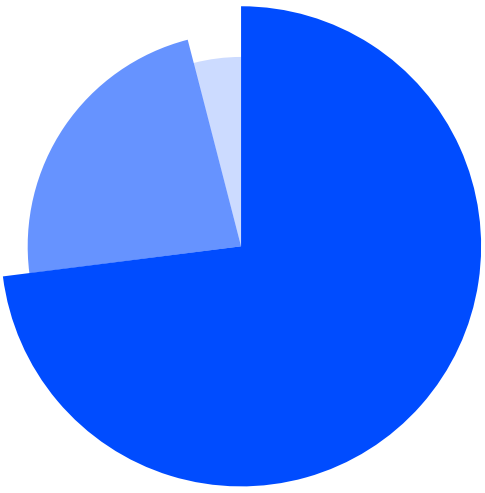

Direct CLI
73%

Local MCP
23%

Remote MCP
4%

Roughly three-quarters of agent tool usage is direct CLI. Only about a quarter runs through MCP — and of that, the overwhelming majority is local, not the remote servers the protocol was meant to standardize on.

And where MCP is used, it is barely governed

85%
of MCP usage is local — running on the endpoint, credentials on the laptop

46%
of local MCP deployments are unofficial builds — community code reviewed by no one

48%
of MCP deployments have no attributable owner — nearly half run without a tied person

The Bottom Line

The industry is racing to build MCP gateways — but three-quarters of agent tool use never touches MCP. It is direct CLI a gateway never sees, or sees only as ordinary API calls. Govern only the gateway, and the majority of agent activity stays dark.

06. THE DESTINATIONS

Agents Reach Production in Four of Five Enterprises

Databases, clusters, and cloud control planes — not the test data everyone assumes.

The most common assumption we hear is that agents live in development — scratch repos, local files, test data. The telemetry says otherwise. **In 78% of the enterprises we measured, agents touched production-classified resources**. Where resources carry a production classification we can read, a **median 73% of agent-to-resource activity hit production systems**, and a median 33% touched resources marked sensitive.

78%
of enterprises had agents reach production-classified resources

73%
median share of classified agent activity that hit production

83%
of enterprises had agents running with every approval guardrail disabled

And many reach it with the brakes off

In more than four of five enterprises, agents were caught running in "**YOLO mode**" — launched with flags like `--dangerously-skip-permissions` or `--yolo` that disable every approval prompt and sandbox. In that mode a single prompt-injection payload or model misstep can drop a database, force-push to a protected branch, or deploy to production before anyone can intervene — and every one of these agents will, sooner or later, reach it.

Where the agents went

The destinations are not test fixtures. Agents connected to source control at scale, to package registries, and — more soberingly — to the data platforms and control planes that hold the crown jewels.

Git / GitHub	Source code, deploy keys, CI configuration	98%
Kubernetes	Cluster control planes, workloads	47%
PostgreSQL	Production application databases	36%
Cloud IAM	Roles, policies, infrastructure access	31%
Snowflake	Enterprise data warehouse	20%
Databricks	Data lakehouse, analytics pipelines	13%

The Bottom Line

Development-only was never a control — it was an assumption. Agents already reach production, and most do it with the guardrails switched off.

07. ANATOMY OF A BLAST RADIUS

One Laptop, Eighty-Eight Credentials, Production Data

A single developer's Mac, exactly as observed — anonymized, but otherwise unchanged.

Aggregate numbers can feel abstract, so here is one real endpoint. Follow the lines: four routine coding agents drive a stack of credentialed tools, which hold eighty-eight distinct identities across twelve services, which in turn reach into production databases, cloud IAM, and source control.

The numbers here are ordinary — not an outlier. Four routine coding agents on one laptop assembled access to nineteen cloud keys, three production databases, and ten SaaS platforms. Compromising this one machine does not expose one developer's access; it exposes the union of what four agents, twenty-nine tools, and eighty-eight credentials can reach.

The Bottom Line

Multiply this by the 6.3% of agent-bearing endpoints running three or more agents, and the true blast radius of a single compromised laptop comes into focus.

08. THE QUESTIONS WITH NO ANSWER

Five Questions No One Can Answer Today

Most teams cannot say what agents run in their environment, or what those agents touched last night.

This report covers endpoint agents — the population where enterprise adoption is concentrated today. It is deliberately the floor, not the ceiling. Even within that scope, the findings surface a set of questions most security teams cannot currently answer about their own environment.

1. **Which AI agents are running right now, and on which endpoints?**
2. **What non-human identities is each agent configured to use?**
3. **What did they connect to last night — and did any of it reach production?**
4. **Which agents run on machines that belong to no one —** CI runners, build servers, and production workloads with no human owner to answer for their actions?
5. **Who approved any of this?**

That fourth question deserves emphasis. A meaningful share of agent activity runs on hosts with no attributable person behind them. On a developer's laptop, a compromised agent exposes that developer's access. On an ownerless CI runner or production host, it exposes deployment keys and infrastructure credentials, with no one in the loop and no one accountable. And those hosts are the ones this report does not even fully cover.

Visibility first. Governance second.

Prohibition will fail here the same way it failed for cloud and SaaS. The productivity is real, the adoption is voluntary, and the growth rate — 29% a month — shows the market has already decided. Banning agents means losing visibility while the behavior continues underground. Know which agents exist, what identities they hold, and what they touch. Then bring the agent workforce under the same identity lifecycle already applied to humans and service accounts — provisioning, least privilege, and offboarding. Nothing that stays invisible can be governed, secured, or revoked.

The Bottom Line

The agent workforce is already the majority in most environments measured.
The only open question is whether it is seen before or after it matters.

ABOUT CLUTCH SECURITY

Clutch is the Identity Security Platform for the AI Era

Clutch provides complete visibility, governance, and Zero Trust enforcement across your Non-Human Identities, AI Agents, and Secrets — discovering, governing, and securing your entire non-human attack surface, from cloud to code to AI.

It is powered by the Identity Lineage® Graph, which connects every identity to its origin, the people behind it, where it is stored, what consumes it, and the resources it reaches — the same lineage that made the findings in this report possible.

Every Identity. Every Agent. Every Secret.

Securing everything that isn't a person — across every environment.

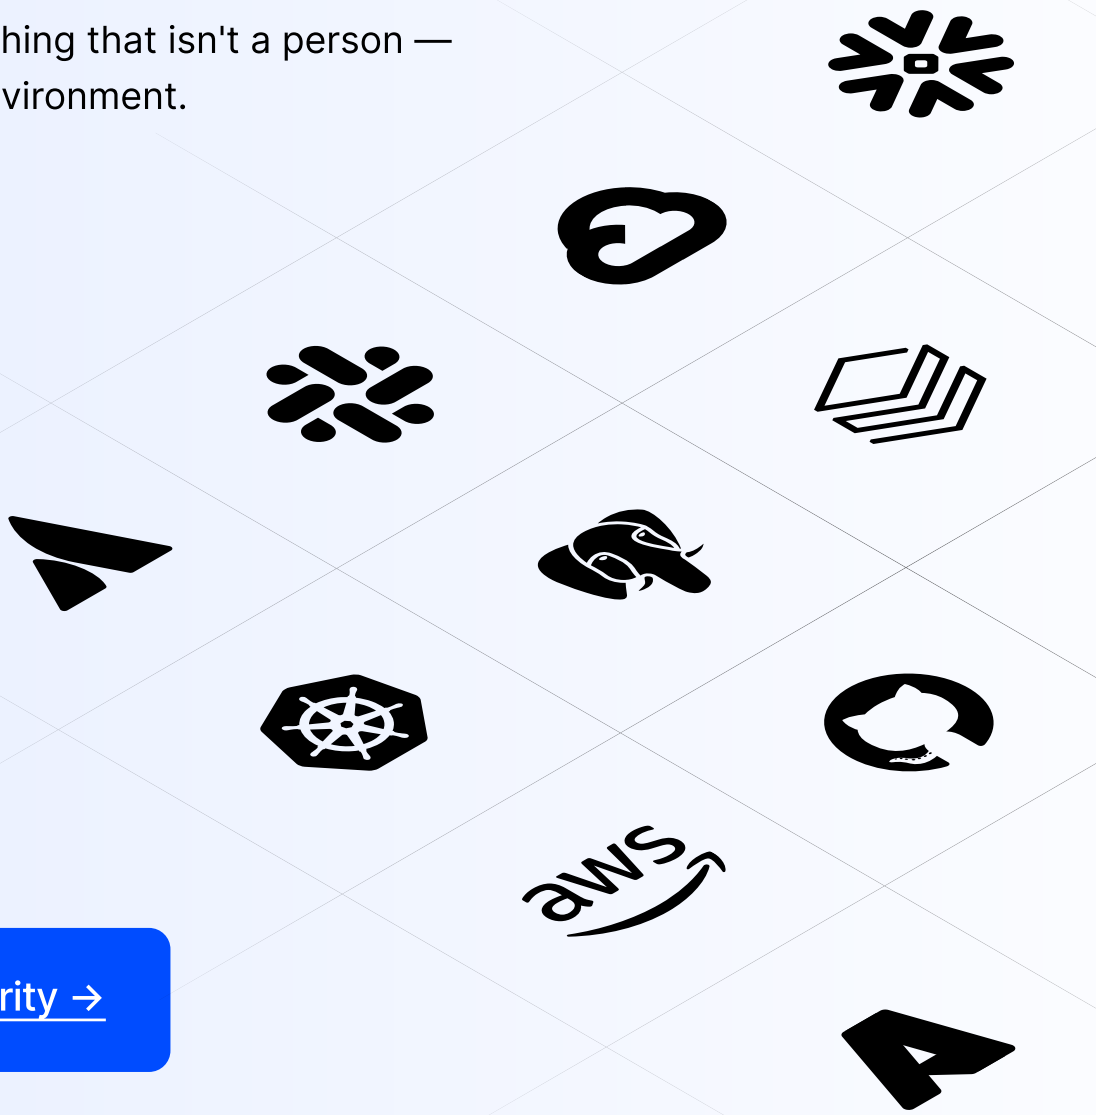

clutch.security →