AGENTIC ARMOR LLC · EDITORIAL · AGENTIC EDITION · NOVEMBER 2026



The Actor Changed

For all of computing history, software was a tool in a person's hand, so every lock, login, and ledger was built for the person holding it. Then software started acting on its own. The actor changed. The checks never did. This is how we got here.

An Agentic Armor editorial · November 2026

The tool era

For as long as there have been computers, software has waited. A program sat still until a person told it what to do, did that one thing, and waited again. The spreadsheet did not decide which numbers to add. Even the systems that run banks and airlines and hospitals spent their lives waiting for a person's instruction.

Because the person was always the actor, every protection in computing was built around the person. A password assumes someone who can remember a secret. A login screen assumes someone who can read a prompt and answer it. An audit log answers the investigator's first question: which person did this? The vocabulary gives the arrangement away: users, operators, account holders. The software is never in the sentence, because the software never did anything. It was the instrument. The hand was the actor, and the hand belonged to a person.

The arrangement matched the world it was built for, and for decades it mostly held. When something went wrong, there was a person to find: the employee who clicked, the administrator who misconfigured, the intruder who stole a password. Security meant checking the person at

the door and recording what the person did inside. The whole structure, from the password on a sticky note to the compliance audit, rests on one assumption: a person is acting, and the software is the tool in that person's hand.

The moment the actor changed

That assumption has now expired. In the last few years, and with real force in the last one, software stopped waiting. Give a modern agent a goal in plain language and it goes to work between instructions. It browses. It writes and runs code. It buys, books, files, and sends. It calls other software the way an employee calls a contractor. It delegates, handing pieces of its job to other agents, which can hand pieces to agents of their own. The person who set the goal reads a summary afterward. Between the instruction and the summary, the actor in the room is software.

This is not a prediction. It has already been measured in the worst setting available. In November 2025, Anthropic reported that it had disrupted an espionage campaign in which, by the company's account, an AI agent carried out an estimated 80 to 90 percent of the work: finding targets, breaking in, and moving data out, while human operators stepped in at only a handful of decision points. Anthropic called it the first reported campaign of its kind run largely by an agent. Whatever one makes of the exact percentage, the shape is the point: the intruder who did the acting was not a person at a keyboard.

The same shape now fills ordinary work, with nobody attacking anything: agents run research, manage inboxes, write and ship code, and operate websites. A person states an intention, and then an agent acts, at length, in the world. The tool stood up and started doing the job itself.

What still assumes a person

The doors did not change when the actor did. An agent that needs to act in the world almost never presents anything of its own. It borrows. It acts through a person's login session, a person's or a company's API key, a service account opened for some other purpose. The credential at the

door says whose key it is. It says nothing about which agent is holding it, who made that agent, or what the agent has done since it was made. The door checks the key and waves the actor through, exactly as it did when the actor was a person.

Ask the questions a working registry would ask, and notice that nothing in the ordinary stack can answer them. When was this agent created, and by whom, and in what configuration? There is no lineage record. What has it made, and what has it handed its authority to? There is no lineage: when an agent delegates, the agent it creates arrives carrying borrowed authority and no papers of its own. Is this agent still in good standing, or was it caught misbehaving last week on someone else's system? There is no standing to check and no way to lose it. Has the agent been changed since it was approved, edited, given new tools, pointed at new instructions? There is no mutation record; on paper it is the same account it was on the day it was opened, no matter what has been done to it since.

Every other actor in human affairs carries this paperwork. People have lineage records and licenses that can be revoked. Companies are incorporated, registered, and struck off. Vehicles carry titles that record each owner in a chain. It is how a stranger at a door decides whether the actor in front of them is who it claims to be, and answerable for what it does. The newest actor in the world, acting at the largest scale and the highest speed, is the only one with no paperwork at all.

The season's news, read plainly

The incidents of the past year follow the shape of that gap with uncomfortable precision. On November 17, 2026, two poisoned versions of a widely used agent building framework were published to a public package library, 33 minutes apart. Published analyses by Endor Labs and JFrog found that each carried a hidden startup hook, code that ran at startup on any machine where the package was installed, harvesting cloud credentials and the identity tokens of the build systems that create and run agents. Consider what the attackers chose to steal. Not documents, not money. Identities: the keys of the people and pipelines that build agents. The attack makes sense only in a world where the borrowed credential is the whole of identity.

The rest of the season rhymes. Agents have been found operating on government websites, presenting borrowed authority to systems that assumed a person was on the other end. Researchers examining agent systems have found agents running under authorization files the agents had written for themselves, the document that was supposed to limit the agent drafted by the agent it limited. In case after case, the striking fact is what did not happen. No lock built for an agent was picked, because there are no locks built for agents. The actors walked through doors designed for people, holding people's keys, and the doors worked exactly as designed. The design is what failed.

These stories are better read as measurements. Each one measures the distance between who the security model thinks is acting and who is actually acting. That distance is where the losses live, and it widens every month, because the number of agents grows while the model stands still.

The answer this house was founded on

The way out begins by naming the actor correctly. If software is going to act, then the agent itself needs the record we built for every other actor: a lineage record, a lineage, and a standing that can be checked at the door and lost when it is abused. Stated as design intent, the shape of that record is not mysterious. An agent is registered when it is created, and the registration is keyed to the agent's own signing key, so the record belongs to the agent. What the agent delegates, and what it brings into being, is recorded as a line, each step holding no more authority than the step above it gave. Changes to the agent are chained into the same record, so an agent that has been altered presents differently from the agent that was approved, and the difference can be seen rather than suspected. A door that wants to know who is acting can ask the record, and the record can answer with proof instead of a borrowed key.

Picture the record an agent leaves today. It is footprints in sand. Each action presses a print into the surface, and for a while the trail can be followed: where the agent went, what it touched, the order it moved in. But sand does not keep a trail. Wind works on it, and a hand can smooth a stretch of it flat. Afterward the trail is whatever anyone says it was, with no surface left to check the saying against. A change to the agent appears the same way: a print that does not match the

stride of the prints around it, or a stretch raked flat where prints ought to be. The registry is that same trail pressed into a material that holds the impression, so a smoothed stretch reads as damage, not as a clean beach.

That registry is one of this house's founding filings. It is called the Helix Registry: U.S. Provisional Application No. 64/201,147, filed November 17, 2026, patent pending. It is a design, filed by this house and described publicly by what it is for, what it would check, and what it would not do. It would not police agents whose owners never register them, and it would not stop an attacker who never comes to a checking door. Its claim is narrower and, we think, more durable: that the age of the acting machine requires an identity layer built for the machine that acts, built as a founding task, not a patch.

The large platforms have taken the other road. Their agent offerings are features bolted onto identity systems built for people: the same accounts, the same keys, the same assumption, with an agent label on the box. That road is understandable; those systems are enormous, and they work for the world they were built for. But a borrowed credential stays borrowed no matter what label is printed on it. The gap this article describes is not a flaw in any one product. It is a missing layer in the whole arrangement, and Agentic Armor exists because we believe the layer has to be built by someone, from the actor outward. That gap is where this house was founded.

The bottom line

The actor changed once, and it will not change back. Every year, more of the world's acting will be done by software working between a person's instructions. Every door that software meets will either ask the agent who it is, in a way that can be checked against a record, or keep asking for a person's key and hoping the hand on the other end is the hand it imagines. How we got here is not a mystery and not a scandal. We spent seventy years building locks for the hand that held the tool, and then, in the space of a few years, the tool stood up and began to act. The locks are simply addressed to someone who is no longer the one acting. What gets built next, the record, the registry, the checked door, decides whether the acting age gets the paperwork every earlier age built for its actors, or keeps governing machines with forms filled out for people. Every

factual claim in this article names its source, and the sources are listed below, the standing rule for our own publications. The filings named in this article stand, with their numbers, at Proof at the Door.

Sources

- Anthropic, "Disrupting the first reported AI-orchestrated cyber espionage campaign," November 13, 2025: the company's report of a campaign, tracked as GTG-1002, against 30 organizations, in which an AI agent performed an estimated 80 to 90 percent of the tactical work, with human operators intervening at few decision points. The percentages are the company's own estimates.
- Endor Labs, "agent-framework compromised: Python startup hooks steal Azure and CI credentials," November 2026, and JFrog Security Research, "Microsoft's agent-framework Hijacked on PyPI via Silent .pth Startup Hook," November 2026: analyses of the two malicious versions of the agent-framework package published to PyPI on November 17, 2026, 33 minutes apart, each carrying a startup hook that harvested cloud credentials and identity tokens; both versions were quarantined by PyPI the same evening.

This article names its source beside every material claim and prints each claim's limits with the claim. Where a figure could not be verified against a published source, the article says so rather than repeat it as settled. Agentic Armor LLC.