

Nuggets Brings Verifiable Authority Enforcement to LangChain and LangGraph

As enterprises move AI agents into production on real actions, langchain-nuggets enforces what each agent is allowed to do at the moment it acts, and proves it.



LONDON, UNITED KINGDOM, July 1, 2026 /EINPresswire.com/ -- [Nuggets](#), the trust infrastructure layer for autonomous AI, today released [langchain-nuggets](#), an open source Python and TypeScript package that enforces authority on every LangChain and LangGraph agent action at the moment it happens.

Access controls determine which systems an agent can reach. They say nothing about whether a specific action, for a specific user, under a specific delegation, right now, is one the agent is authorised to take. That question goes unanswered in most production agent stacks, and for regulated industries deploying agents that move money, change records, or act on behalf of customers and partners, that is no longer acceptable.

langchain-nuggets closes that. It intercepts every tool call before it executes, checks five constraints per call against a signed delegation credential, and either allows the action and emits a cryptographic proof artifact, or blocks the call entirely. Revocation takes effect immediately. Consent is enforced the same way, withdraw it, revoke the delegation, and the agent is stopped at the next call.

The proof artifact is what distinguishes langchain-nuggets from existing authorisation and guardrail layers. Every allowed action produces a signed Action Receipt: who acted, on whose authority, under which constraints, with what declared intent. The receipt references hashes of parameters and results, not the raw data, making it auditable without being a surveillance record. In production, receipts are signed with RS256 against Nuggets' published public key and verify independently, by a partner, a regulator, or a system on a different cloud, without access to the originating organisation's systems or logs.

Trust travels with the agent. That holds across companies, vendors, and clouds, and traces back

through a verified identity chain to the human who ultimately authorised the action.

"Access decides what systems your agent can reach," said Seema Khinda Johnson, Co-founder & CCO of Nuggets. "Authority decides whether this specific action, for this user, right now, is one it's allowed to take, and leaves proof it was. That's what langchain-nuggets gives you, and it's what regulated enterprises need before they can put agents on the actions that matter."

langchain-nuggets 1.0.0 is available now on PyPI (pip install langchain-nuggets) and npm (npm install @nuggetslife/langchain-nuggets). The cross-org authority demo runs entirely locally without a Nuggets account. Full documentation and source code are at github.com/NuggetsLtd/langchain-nuggets.

About Nuggets

Nuggets is the trust infrastructure layer for autonomous AI, built to govern what agents actually do at the point of execution. Its platform enforces identity, authority, policy, intent, consent, and accountability at the moment an action is taken, producing cryptographic proof across cloud and on-premise systems. Nuggets is headquartered in London and operates across North America and Europe.

Dom Gilmore

Nuggets

contact@nuggets.life

Visit us on social media:

[LinkedIn](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/923299069>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.