Reimagining the Root of Trust

A Hardware Foundation for Deploying AI Agents Without Accidents

Enclawed LLC | Tracy, California, USA | **Patent Pending**

The problem

Autonomous AI agents now take real, often irreversible actions, and most agent runtimes do nothing to stop a hijacked agent or prove what happened. Strong software security—like Enclawed's—already closes those gaps today and is the essential foundation; on its own it delivers protection that off-the-shelf agent runtimes simply do not. For the most sensitive deployments, organizations also want assurance that holds even against a compromised host or a physical attacker—a guarantee software alone cannot make for anyone.

The idea

In addition to Enclawed's software security—not as a replacement for it—we are building a **tamper-proof hardware module**: a self-contained "root of trust" that verifies and authorizes every sensitive action *inside* a sealed, shielded boundary, with cryptographic keys that **never leave the device**. Instead of one verifier, it uses **several independent verifiers that must agree**, and it keeps an on-device, tamper-evident record of everything it approves. It is an optional hardware layer for the highest-assurance environments; the software stands on its own without it.

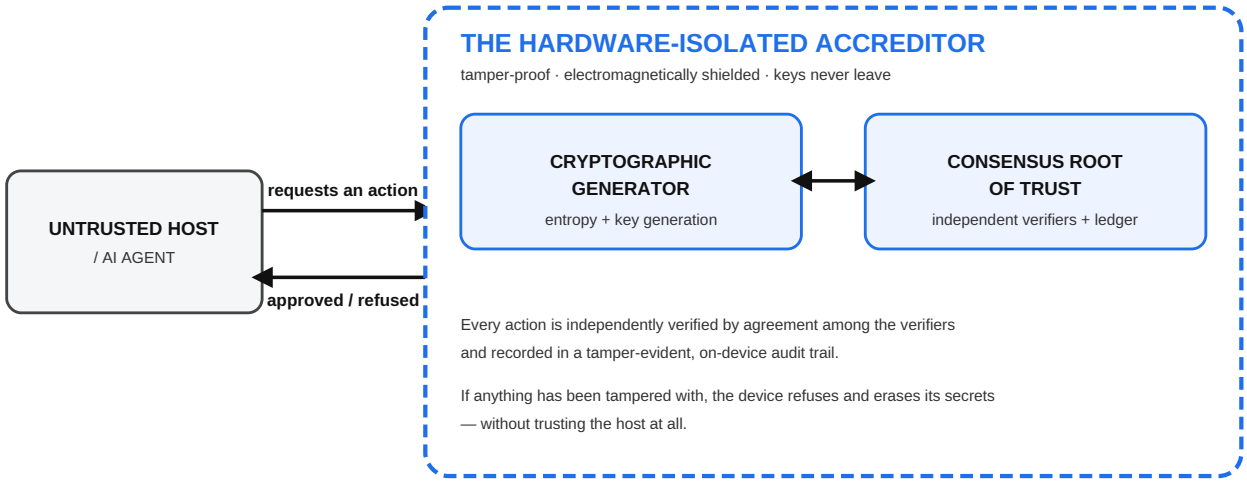

Why it is different

- **No single point of failure.** The secret key is split across independent verifiers; cracking one reveals nothing. Only a quorum acting together can act.
- **Trust by agreement, not by one check.** Defeating it would require subverting many isolated verifiers at once — inside a sealed enclosure.
- **It watches the software while it runs**, not just at startup, and refuses the moment something is altered.

- **It does not trust the host.** Its protections work even if the surrounding computer is fully compromised.
- **Resists physical and electromagnetic attacks**, erasing its secrets if tampering is detected.
- **Designed to meet or exceed FIPS 140-3 Level 4** — the highest physical-security tier — with protections the standard does not require.

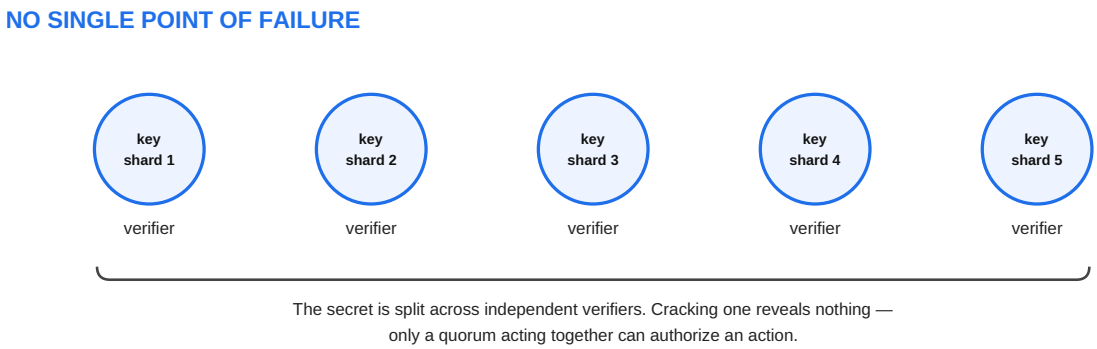

Where it fits

Wherever software must be trusted to act on its own and a mistake is costly: autonomous AI agents first, and more broadly any high-assurance software service that needs stronger protection than today's security hardware provides.

Status & contact. Patent pending (U.S. provisional application filed 2026). The technology is available for licensing and partnership. Alfredo Metere, Founder & CEO, Enclawed LLC — alfredo.metere@enclawed.com.