



Command, control and agentic software, and private AI models, with a governance plane that trusts no agent by default.

Moonwalker builds command, control and communications software, and agentic software for any mission, and the governance plane they can run on. The plane applies zero trust to the AI itself. **Every agent action is checked against policy where it happens**, at the file, the command, the browser and the API. An action beyond granted authority is denied, sensitive actions wait for a human, and every action is recorded outside the agent in a tamper-evident record. It installs air-gapped or connected, on Windows and Linux, and **the customer keeps full data sovereignty**, with your AI provider, your storage and your keys.

QUALIFICATIONS

- SBA-certified Service-Disabled Veteran-Owned Small Business
- Nontraditional defense contractor, no cost share on a prototype OT
- Commercial term license, no metering, no SaaS dependency

ZERO TRUST ROLES FOR AI AGENTS

DECIDE	ENFORCE	ACCESS	RECORD	WATCH
Policy decision point for every agent action	Enforcement at the application, API, data and endpoint	Least privilege and privileged-activity control for agents	Data integrity, audit and compliance evidence	Monitoring of agent actions and central policy management

CORE COMPETENCIES

Policy decision and enforcement	Agent access and privileged activity	Evidence, audit and data integrity	Design and custom build	Private AI models on your own hardware
■ Policy held as data, outside the model ■ Checked at the file, command, browser, desktop and API ■ Denied beyond granted authority, recorded either way ■ Risk-scored actions with a hard stop ■ Fail-closed, air-gapped or connected ■ Governed browser and sealed web capture	■ Least-privilege access for non-person identities ■ Just-enough access by tool, path and action class ■ Human approval workflows for elevated actions ■ Agent session management ■ Any agent, any harness, any AI provider	■ Tamper-evident record written outside the agent ■ Hash-chained entries, signed checkpoints, SHA-384 ■ Customer-held storage and keys ■ Human-authorized access, separation of duties ■ Compliance-ready evidence, incident and after-action reconstruction	■ Discovery, requirements and AI systems architecture, with or without the build ■ Systems engineering and technical assistance (SETA) ■ Command, control and communications (C3) software ■ Custom agentic solutions for enterprise and government ■ Governed offensive-security and research tooling, built to order ■ Multi-agent coordination and Model Context Protocol connectors	■ Language models built from scratch on your own data, air-gapped, weights you own ■ Training that learns from machine-checked results alone; the model never grades itself ■ Task-specific helpers built on site for isolated networks, proven against a locked test before use ■ Leak-checked evaluation, test rules written before the run ■ Digital twins, built to order

DIFFERENTIATORS

- **Zero trust applied to the AI itself.** Network and identity zero trust stop at the login. Moonwalker keeps checking after it, on every action the agent takes.
- **Enforcement where the action happens.** An agent cannot route around a control that lives at the surface it touches.
- **Data sovereignty.** The record is written to storage you select under keys you hold, inside your own boundary, so neither an agent nor the vendor can open it.
- **Lab neutral.** No model of our own attached to the plane, no data brokered, no share of model usage. You pick the AI you trust, including a private model we build for you, with weights you own.
- **No authorization of its own to wait for.** The software runs inside your boundary under your authorization and inherits the enclave's impact level. Software bill of materials, hardening guide and control responsibility matrix ship with the license.

KEY PERSONNEL EXPERIENCE

Dr. Clinton P. Cuzzort

FOUNDER, CHIEF EXECUTIVE OFFICER AND LEAD ENGINEER

- Designed and built the Moonwalker governance plane, in production use in two organizations on separate hardware
- Designed AI systems from research to build-ready design, including agent runtimes, multi-agent coordination and resource allocation, source-traced document engines, post-mission debrief and re-planning, agent memory, sealed records, agent identity and credential brokering, recovery and shutdown, voice interfaces, and drone and robot governance
- Full-stack AI development from kernel and binary layers to multilayered cognitive memory fabrics
- Security and compliance frameworks that control autonomous agents in enterprise environments
- Evaluation harnesses with fixed inputs, pinned configurations and a stated line between system failure and model behavior
- U.S. Army radar and missile electronics, 2004 to 2008, deployed to Iraq
- Industrial controls, PLC logic and HMI design, Nucor Steel
- Strategic planning and internal controls, Overseas Private Investment Corporation
- Doctor of Medicine, clinical training in high-consequence decision environments

DEMONSTRATED CAPABILITY, LIVE OR BY RECORDED WALKTHROUGH ON REQUEST

LIVE GOVERNANCE	BOUNDARY ENFORCEMENT	GOVERNED OFFENSIVE WORK
Normal work executing through the control plane, with the audit record populating in real time.	An action outside its permissions denied at the boundary, with the denial recorded and risk scored.	Offensive security and research tooling run through the control plane like any other agent, with a sealed record of everything it did, proven live against the company's own boundary from a separate attack system.

ENTITY	UEI	CAGE	NAICS	PSC	ACCESS
Moonwalker Inc. Mississippi corporation, small business, SDVOSB	W5TVGMMNGWU3 SAM Active, All Awards	23QR6 Section 889 represented	513210 primary 541512 · 541519 · 541715	7J20 · DJ10 · DA10 Software, IT and AI services	Direct award, OT, sub Supported 90-day evaluation under micro- purchase