Castellan

The Agentic Chassis — AI agents that act with money and authority — inside hard limits, on a complete audit trail.

> **THE NEED**
>
> Every office is being told to adopt AI, but an agent that can spend funds, approve actions, or commit an organization cannot be a black box. Program offices need automation they can bound, audit, and answer for — software that acts, and accounts for every action. That accountability layer is what we build, and it is already running in production.

WHAT IT IS AND WHAT IT DOES

Castellan is production infrastructure for **governed AI agents**: each agent operates only inside a **user-defined mandate** (scope, dollar amount, destination), moves money through **escrow and verified payouts** on regulated U.S. rails, and writes **every decision to an audit log** with its expected and actual outcome. Anything outside the mandate stops and asks a human. One backbone serves every deployment, so a new mission workflow inherits the guardrails — identity verification, spend limits, injection defenses, decision logging — on day one. Each deployment runs its own named Castellan agent: Don Aduana handles customs compliance for DirectoPaRD, and USALink runs its domestic purchasing agent on the same backbone.

EXPERIENCE

DirectoPaRD — cross-border commerce (U.S. ↔ Dominican Republic)
Developed and operate an AI purchasing and customs-compliance portal: the agent prepares and executes purchases inside user-set mandates, screens customs and export rules before money moves, and settles through escrow with verified sellers. Value delivered: compliance screening and purchase preparation that previously required manual broker work now run automatically, with a human approving anything outside mandate.

USALink — Castellan operating domestically
Employs Castellan for U.S. domestic agentic commerce: mandate-bound purchasing with full decision logging on American rails. Solves the gap between consumer AI assistants and accountable procurement — every automated purchase is bounded, logged, and reviewable.

Built and operated in-house since 2025 on one shared codebase.

WHAT IT DOES FOR A PROGRAM OFFICE

- **Governed micro-purchasing:** agents execute routine buys under hard spend limits with a complete audit file for every action.
- **Automated rule checking:** customs-style compliance screening with a human-in-the-loop for exceptions.
- **Verified onboarding & payouts** for distributed field personnel and vendors.
- **Audit-ready autonomy:** expected vs. actual outcome recorded per decision — no black box to defend.

DIFFERENTIATORS

- **Mandate-bound by design** — outside the mandate, the agent stops and asks.
- **Hardened against prompt injection** — layered defenses verified in production.
- **100% U.S.-built AI stack** — U.S. models, cloud, and payment rails; no foreign model providers.
- **User owns the data** — never sold, never shared.

COMPANY DATA

UEI	C4KGWJ6T4BB3
CAGE	22QH8
NAICS	541715
SBC ID	SBC_002680767

CERTIFICATIONS & STATUS

SAM.gov Active · Florida LLC
Small Business · Minority-owned
CMMC Level 2 self-assessment in progress (NIST SP 800-171)
Registered: DSIP, Grants.gov, Research.gov, SBA Company Registry

TECHNICAL STACK

TypeScript / Next.js · PostgreSQL · Stripe Connect & Treasury rails · Anthropic Claude (U.S.) · Vercel (U.S.) · CI policy guard enforcing code doctrines before every deploy

Maciel DeJesus, CEO · 267-232-7614 · maciel@condestgermaintech.com
1717 County Road 220, Ste 1605, Fleming Island, FL 32003 · condestgermaintech.com