



# WATCHFIRE AI

*Human agency, engineered.*

CAPABILITY STATEMENT CAGE 1AET1 UEI EJWWEAJP3ML7 SDVOSB Application Pending

Watchfire AI builds production AI systems where human approval gates block execution at consequential points, the audit trail is generated automatically as the system runs, and **oversight is enforced in the architecture** — not added as a governance layer on top.

## CORE COMPETENCIES

- **Auditable AI for high-stakes missions.** Architecture-first design where every consequential action requires meaningful human approval and produces traceable evidence.
- **Adversarial verification.** Multi-tier pipelines that decompose AI outputs into material claims and produce defensible artifacts with per-claim verdicts.
- **Probabilistic decision support.** Structured forecasting with stage-gated human approval, replacing gut-feel analysis with calibrated, defensible forecasts.
- **Federal-grade AWS deployment.** Production AI systems running on AWS today, with appropriate controls for sensitive and mission-critical use.

## DIFFERENTIATORS

- **TS/SCI/FSP cleared SMEs.** Active clearance enables work in cleared environments without processing delay.
- **Architecture-first oversight.** Approval gates and audit trails implemented in code, not policy theater. Defensible to technical evaluators.
- **Federal AI delivery pedigree.** Founder experience spans both AI-focused prime (C3 AI) and hyperscaler (AWS) sides of the federal AI stack.
- **AWS-native production today.** Systems running in production now, not roadmap. Custom builds available for on-prem and edge.

## PRODUCTION SYSTEMS

**DECON** Tiered adversarial verification engine. Decomposes input into material claims, runs cross-model adversarial disconfirmation, verifies cited sources against external bibliographic databases, and produces a defensible report with per-claim verdicts and traceable evidence.

**BARRACUDA** Structured probabilistic forecasting for high-stakes decisions. Operationalizes resolution criteria, runs multi-stage evidence gathering, models scenarios with explicit probability allocation, and applies a bias-check synthesis pass to produce calibrated forecasts with full audit trail.

**Additional systems** *BLOODHOUND (coordinated inauthentic behavior detection), MERIDIAN (adaptive learning), Border-Ready (travel entry validation), and four others in production.*

## FOUNDER EXPERIENCE

*Founder past performance, listed for context. Not Watchfire AI past performance.*

**Jack Roelofs** Founder & Nightwatchman. **C3 AI** (Director, Federal Business Development): MDA SHIELD, Army PEO C3N NGC2, HHS SHARE BPA. **AWS National Security** (Defense and Intel AI Transformation Leader): JWCC, WILDANDSTORMY, C2E. **U.S. Military** (Army SOF and Infantry Officer, Air Force Communications Engineer NCO).