

COMPANY PROFILE

Business Type

Service-Disabled Veteran-Owned Small Business (SDVOSB)

SBA VetCert

Approved 2026-04-21

CAGE / UEI

19AK6 / S6DTCH6ULHD1

SAM.gov

Active & current

CMMC Posture

Level 2 (self-assessment) - DFARS 252.204-7012 / 7019 / 7020

Foreign Nationals

None

NAICS Codes

541512, 541519, 518210, 541511, 541715

Website

oubliettesecurity.com | Demo: oubliettesecurity.com/demo

DIFFERENTIATORS

- > Working software, not vaporware - two products shipping on PyPI since Feb 2026; 1,500+ tests passing; live demo at oubliettesecurity.com/demo.
- > SDVOSB - SBA VetCert approved 2026-04-21; qualifies for FAR 19.1406 sole-source and set-aside competitions across DoD, VA, and civilian agencies.
- > Peer-reviewed methodology - UK AI Safety Institute inspect_evals PR #1358 contributed under independent review.
- > Open-source-first, ITAR-aware - permissive licenses; air-gappable Ollama backend; CMMC L2 (self) today, certifiable for Phase II / CUI.

PRODUCT SUITE

Product	Status	Description
Oubliette Shield	PyPI v1.0.0	Runtime AI firewall, 5-stage detection + deception
Oubliette Dungeon	PyPI v1.0.1	Adversarial testing engine, 57 scenarios
Agentic Cyber Defense	Phase I	Four-agent framework, operator-on-the-loop UI

PAST & PENDING FEDERAL ACTIVITY

Effort	Status
DIU MYSTIC DEPOT LOE1 (PROJ00625)	Submitted 2026-03-24
SCO BIAO white paper - BAA HY0233-SCO-24-BAA-0001	Drafted
DARPA I2O abstract - BAA HR001126S0001	Drafted
NATO RFIP-ACT-SACT-26-38 concept paper	Drafted
DLA SBIR DLA26BZ02-NV005	In prep, due 2026-06-24

PROCUREMENT VEHICLES

SDVOSB sole-source (FAR 19.1406) - SDVOSB set-aside (DoD/VA/civilian) - SBIR/STTR Phase II/III - Other Transaction Authority (DIU and similar) - DoD OSBP Mentor-Protege - Subcontracting on prime-led efforts.

CORE COMPETENCIES

AI Security

Runtime detection, cyber deception, and adversarial robustness for production LLM systems. Five-stage pipeline (sanitizer, pre-filter, ML classifier F1=0.98, RAG guard, LLM judge). 12 provider, 9 SDK integrations.

Agentic AI Security

Multi-agent cyber defense with safety-gated autonomous tooling; operator-on-the-loop review. ATT&CK-aligned planning; citation-bound vuln research; CALDERA-emulated substrate.

Adversarial AI Testing & Red Teaming

57-scenario suite under UK AI Safety Institute review (inspect_evals PR #1358). Multi-provider comparison; React dashboard; Garak / PyRIT integration.

Cyber Deception

Honeypot, tarpit, and redirect modes in the AI firewall. STIX 2.1 threat intel as a byproduct; honey tokens alert on attacker reuse.

Cognitive Warfare Validation

Prototype harness for testing cognitive-warfare AI models under safety and audit constraints.