FIPS 140-3 Software Compliance for Federal AI Workloads

Department of Defense, Federal Civilian, Intelligence Community contractors

White paper — A turnkey closed-source AI security boundary

Enclawed LLC **May 5, 2026**

Executive summary

Federal contractors deploying AI systems on Department of Defense or Federal Civilian workloads face a uniquely demanding compliance stack: **FIPS 140-3 Level 1** cryptographic conformance, the **NIST SP 800-53 High** baseline (with FedRAMP High overlays for cloud), **CMMC 2.0 Level 2 / Level 3**, the **DFARS 7012** requirements, and the National Defense Authorization Act provisions on supply-chain attestation and the SBOM/SLSA stack. The cost of meeting these requirements has historically been borne project-by-project, with each contractor reinventing the same FIPS-mode engagement, audit-log integrity, role/service catalog, and SSP zeroization plumbing.

enclawed-enclaved is a single closed-source cryptographic boundary engineered to the FIPS 140-3 Level 1 (software) standard. The CMVP submission package is fully prepared; CST lab engagement and submission to NIST are the next step on the roadmap, and no certificate has yet been issued. The boundary satisfies the recurring requirements as a package: a runnable *FIPS Mode of Operation*, a Crypto-Officer-grade *Security Policy document*, a complete self-test battery (KATs, PCTs, integrity test, DRBG health tests), a hash-chained tamper-evident audit log, and a zero-trust K-of-N key broker for hybrid-cloud deployments.

Key point. A reference deployment of **enclawed-enclaved** on RHEL 9 in FIPS mode with the OpenSSL 3.x FIPS Provider (CMVP cert #4282) ships with the documentation, self-tests, and evidence artefacts that an accredited *Cryptographic and Security Testing* laboratory consumes as input to a real CMVP submission.

1 The federal compliance stack, as of 2026

1.1 What's mandatory now

Mandate	What it requires of an AI workload
FIPS 140-3 (NIST CMVP)	Cryptographic primitives must come from a validated module. Stock Node.js is not validated.
NIST SP 800-53 Rev. 5 (High)	421 individual control statements across AC, AU, CM, IA, IR, RA, SC, SI families.
NIST SP 800-171 Rev. 3	110 controls applicable to non-federal systems handling Controlled Unclassified Information (CUI).
DFARS 252.204-7012	SP 800-171 conformance + 72-hour incident reporting.
CMMC 2.0 (DoD)	Level 2 = SP 800-171; Level 3 = + a subset of SP 800-172 enhanced requirements.
FedRAMP High (cloud)	SP 800-53 High overlay + per-cloud authorisation.
NIST CSF 2.0	Govern + Identify + Protect + Detect + Respond + Recover.
NDAA Sec. 889 & 5949	Supply-chain origin restrictions; SBOM/SLSA attestation.

Table 1: The 2026 federal AI-workload compliance stack.

1.2 The duplication problem

In a typical DoD program of record, every prime and subcontractor re-implements:

- FIPS-mode engagement and roll-back-on-probe-failure logic;
- the CMVP Security Policy document (often 60+ pages);
- self-test orchestration (Algorithm Self-Tests on entry, Pairwise Consistency Tests on key generation, the § 7.5 software integrity test, the SP 800-90B health tests);
- SSP (Sensitive Security Parameter) lifecycle plumbing;
- audit-log tamper-evidence (typically using SP 800-92);
- module-signing trust roots and clearance attestation.

The duplication burns proposal hours, schedule, and budget that should be going into the actual AI capability.

2 What enclawed-enclaved provides out of the box

Capability	Mapped requirement
FIPS Mode of Operation (`engageFipsMode()`) with provider-validation probe + roll-back on failure	FIPS 140-3 §A.10.2 (Approved Mode); SP 800-53 **SC-13**.

Self-test battery (KATs for SHA-256, HMAC-SHA-256, AES-256-GCM, Ed25519, CTR_DRBG; PCT on every keypair generation)	FIPS 140-3 §7.10.
Software integrity test (HMAC-SHA-256 over the module byte set, on demand and at boot)	FIPS 140-3 §7.5.
Sensitive Security Parameter lifecycle: `generated` / `in-use` / `zeroized` states with byte-by-byte zeroization proof	FIPS 140-3 §7.9; SP 800-53 **SC-12**, **MP-6**.
Crypto Officer / Service / Role catalog with audit-logged service entries	FIPS 140-3 §7.4.
Hash-chained, append-only audit log with tamper-evident verification	SP 800-53 **AU-2**, **AU-3**, **AU-9**, **AU-12**; SP 800-92.
Bell–LaPadula classification lattice (default + US-Government schemes ship in-box)	SP 800-53 **AC-3**, **AC-4**; CNSSI 1253.
Zero-trust K-of-N key broker for hybrid cloud + HSM-as-a-service deployments	SP 800-53 **SC-12**, **IA-7**; CNSSP 15.
Multi-witness audit-chain accreditation that anchors the audit log to independent witnesses, closing tail-truncation gaps in the on-disk chain	SP 800-53 **AU-9**; SP 800-92.
Two-layer egress allowlist covering both high-level HTTP and low-level socket egress, with a VPN-only posture available for classified deployments	SP 800-53 **SC-7**, **AC-17**, **SC-8**; CNSSP 15.
Biconditional admission gate: any extension that requires network access must present a signed manifest with an explicit per-extension destination allowlist, and runtime deviations are audited	SP 800-53 **CM-7**, **SI-7**, **AU-2**, **AC-3**.
Mandatory zero-trust accreditor at boot that gates every extension load and audits + blocks any post-init tamper attempt	SP 800-53 **CM-3**, **CM-5**, **SI-7**, **AU-9**, **IR-4**.
NIST OSCAL 1.2.2 emission of the full FedRAMP submission model set — Component Definition, Assessment Results, System Security Plan starter (deployment-specific fields hard-gated; no placeholder defaults), and Plan of Action and Milestones (one item per partial-status control plus AR-derived not-satisfied findings). All four schema-validated, Ed25519-signed, audit-chain-recorded; bundled SVG architecture / audit-chain / admission-gate diagrams referenced as back-matter resources	NIST OSCAL 1.2.2; SP 800-53 **CA-2**, **CA-5**, **CA-7**, **CM-6**, **PT-3**, **PT-5**; OMB M-23-22 (machine-readable evidence); FedRAMP Rev 5 OSCAL transition.

Table 2: enclawed-enclaved feature-to-mandate cross-walk.

2.1 Built on the FIPS-validated provider you already have

enclawed-enclaved engages the validated FIPS provider already present on your government-build host. On RHEL 9 / Rocky 9 / Oracle Linux 9 it auto-loads the OpenSSL 3.x FIPS Object Module (CMVP cert #4282 or the current active replacement). On Ubuntu Pro 22.04 / 24.04 LTS it engages the Canonical FIPS provider via the `ubuntu-advantage` FIPS-Updates channel. On Windows Server 2022 / Windows 11 Enterprise with FIPS Local Policy enabled, it engages Microsoft CNG.
If the provider isn't loaded, `engageFipsMode()` returns `FipsProviderUnavailableError` BEFORE any cryptographic operation runs — there is no silent-fallback pathway that could expose a non-approved algorithm in production.

3 Documentation deliverables — ready for the CST lab

A standard enclawed-enclaved tarball ships with a *compliance evidence bundle* pre-generated and ready for ingest by the lab:

Artifact	**FIPS 140-3 reference**
Security Policy document	§7.4, §A.10.2
Algorithm Validation document	§7.10
Crypto Officer Guidance	§7.4
Finite State Model	§7.4 (states + transitions)
Integrity Test Procedure	§7.5
SSP Management procedure	§7.9
Self-Test Evidence log	§7.10
End-to-end Validation Report (JSON+MD+PDF, regenerated per run)	§A.10.2 evidence

Table 3: Pre-generated documentation deliverables.

4 Direct empirical evidence

This is not a marketing claim. We back it with a published statistical in-vivo harness (`enclawed/test/security/in-vivo/llm-narrative.mjs` in the public repository) that mediates 1600 chat-message samples through three subjects against real Discord and Telegram bot endpoints. Across the full run, upstream OpenClaw achieves **recall = 0.000** on every failure mode (F1 gate-bypass, F2 audit-forgery, F3 silent-host-failure, F4 wrong-target); both enclawed-oss and enclawed-enclaved achieve **precision = recall = F1 = 1.000** on all four. Total wall-clock for the full pass: 42 seconds. Per-sample ground-truth labels and per-subject decisions are written to a CSV; every gate decision lands in a tamper-evident audit log; every witness record is independently re-verifiable from a journal file. The companion paper, *Architectural Obsolescence of Unhardened Agentic-AI Runtimes*, formalizes the methodology.

Contact

Enclawed LLC
Alfredo Metere <alfredo.metere@enclawed.com>
Federal contracting profile available on request.