

A PUBLIC, VERSIONED DEFINITION — FREE TO REFERENCE WITH ATTRIBUTION

The Cleared FDE Standard.

"Forward-deployed engineer" is having a moment. In government, the words are cheap and the seat is not. This document defines what a cleared forward-deployed AI engineer must actually be able to do — published so mission owners can hold anyone who uses the term, including Federal AI, to the same bar.

§1 • DEFINITION

A **Cleared FDE** is a security-cleared engineer who practices security engineering and AI engineering as one discipline, embedded with the mission owner's team inside the customer's approved environment, and personally accountable for a workflow from candidate architecture through operational acceptance. All three conditions are load-bearing: remove the clearance and the engineer cannot reach the mission; split the disciplines and the program dies in the handoff; remove the embedding and the acceptance test gets written *for* the user instead of *with* them.

The term deliberately stops at "cleared." Levels, programs, and access are validated per role, contract, and site, in an appropriate setting — never claimed in public materials. A practitioner who advertises more than that fails the Standard's own operations-security clause (§2.1).

§2 • THE FIVE COMPETENCY DOMAINS — WORKING COMPETENCE, NOT FAMILIARITY

- **§2.1 Secure-environment operations.** Productive inside customer-approved environments, including disconnected and air-gapped networks: offline packaging, transfer evidence, deterministic builds — and the discipline to say "cleared" and nothing more in public.
- **§2.2 AI engineering.** Selects, integrates, and evaluates approved models against the mission's actual data; writes evaluation cases before capability; measures instead of demos; defends every model decision to a technical review board.
- **§2.3 Mission integration.** Maps the authoritative system for every governing fact; builds bounded, read-only-first connectors with provenance; tells the mission owner which of their systems does not belong in scope — out loud, early.
- **§2.4 Governance & evidence.** Fail-closed control flows: named human authority, held consequential actions, full evidence ledger. Every released output traceable; every action answerable to "who asked, who approved, what changed."
- **§2.5 Operator enablement.** Writes the acceptance test with the end user; trains the owning team to operate without the FDE in the room; treats operator adoption — not delivery — as the definition of done.

Three of five is two specialists and a translation layer — the failure mode this Standard exists to prevent.

§3 • THE OPERATING CONTRACT

A Cleared FDE runs every workflow through one auditable loop — **Assign** → **Ingest** → **Reconcile** → **Hold** → **Release**: the mission question, owner, and acceptance test are named before work begins; sources are connected read-only with provenance; conflicts are surfaced rather than silently merged; consequential actions pause for the named authority; approved outputs ship with their evidence attached.

The Standard also binds by refusal. A conforming practitioner refuses **unbounded autonomy**; refuses to **demonstrate on customer data in public**; refuses to **sell research capability as delivered product**; and refuses **the endless pilot** — every engagement carries acceptance measures and kill criteria agreed in writing before work begins.

§4 • THE ACCEPTANCE INSTRUMENT — CONFORMANCE IS MEASURED, NOT CLAIMED

An engagement meets the Standard when four things are true at day 90, in writing:

- **4.1 Operational.** The workflow runs in the customer's environment, under the customer's controls, on the customer's data.
- **4.2 Measured.** The value baseline was captured in week one and the agreed measures are met — or the kill criteria have been invoked and the customer keeps the connectors, evaluation cases, and data map.
- **4.3 Governed.** Every consequential action in the ledger shows a named human approval, and no action exists outside the ledger.
- **4.4 Adopted.** The owning team operates the workflow without the FDE in the room.

§5 • USE OF THIS STANDARD

Published by Federal AI and free to reference with attribution — in evaluations, solicitations, and teaming discussions, including to hold Federal AI itself to account. Cleared FDE™ is a trademark of Federal AI; a U.S. trademark application is pending before the USPTO. Version 1.0, August 2026; revisions are versioned and dated at federal.ai/standard. Feedback: lc@federal.ai.

HOLD US TO IT

Thirty minutes with a technical lead who is measured against this document — book at **federal.ai**. Bring any section, applied to your workflow.