[CAPABILITY PROFILE · AUTHREX-05]

FLAME · FORCED LATENCY FOR AUTHORITY-MANAGED ESCALATION

Engineer the time between a proposed irreversible action and its execution so that a human veto is exercisable before the outcome, with abort as the only default on timeout.

MODELED **TESTED**

Name variants: Patent filing title: Flash War Latency Architecture (preserved as the source document name). Same framework; deposit v5.11.

PRIMARY FUNCTION	Mandatory deliberation window before an irreversible action, sized by a dynamic delay function D(A, tier, domain) over authority, consequence tier, and domain; five-state circuit breaker with cryptographically signed transitions and a keep-alive heartbeat; timeout defaults to abort, never execute.
MISSION AREAS	Engagement authorization chains; time-critical command and control; any decision with an irreversible effect, including counter-UAS mitigation and automated control writes in operational technology.
EVIDENCE ANCHORS	Zenodo DOI 10.5281/zenodo.19015618; seeded browser simulation; formal-coverage record: FLAME.tla surrogate countermodel with a published strong-form counterexample and a stated requirement.
CURRENT STANDING	Demonstrated in simulation; the deliberation mechanism is specified and simulated; the formal record shows the naive clock property can be satisfied with zero real time elapsed; not independently reproduced.
SELF-ASSESSED TRL	TRL 3. No formal TRL determination has been performed.
INTELLECTUAL PROPERTY	U.S. provisional application 64/005,607, filed March 14, 2026 (receipt 74858888), reported as filed; unexamined; confers no enforceable rights.

[MISSION PROBLEM]

A veto that cannot be exercised before an irreversible outcome is theoretical. Machine decision latency and human deliberation time are usually engineered separately, so a system can be formally supervised and practically unsupervisable. The failure is compounded when the fallback on a lost confirmation is to proceed. The consequence is an irreversible action executed inside a window no human could have used.

[OPERATIONAL RELEVANCE]

- Engagement authorization chains for uncrewed systems, where the confirm window must scale with consequence.
- Counter-UAS mitigation timing, where the decision window is seconds and the effect is irreversible.
- Operational technology: mandatory deliberation before automated load-shedding or control writes (BLADE-INFRA-OT, ICS-GATE).
- Spacecraft operations under light delay, where evidence age and budget gate the action (DELAY then ABORT if not refreshed).

[CAPABILITY DESCRIPTION]

FLAME treats latency as an engineered authority parameter. Before any irreversible action it imposes a deliberation window sized by D(A, tier, domain): higher consequence buys the human more time by construction. A five-state circuit breaker with cryptographically signed transitions enforces the window; a keep-alive heartbeat detects a stalled supervisor. If the window expires without the required confirmation the outcome is abort, not execute. Where mission policy explicitly authorizes autonomous execution after a bounded hold, that authority must already exist in the pre-authorized envelope; timeout never creates it. The formal record adds a requirement learned from a counterexample: the deliberation clock must advance only by elapsed time, because a clock that can be advanced by bookkeeping alone can be satisfied with zero real time elapsed.

[FLAME · ARCHITECTURE, ASSURANCE, EVIDENCE]

[TECHNICAL ARCHITECTURE]

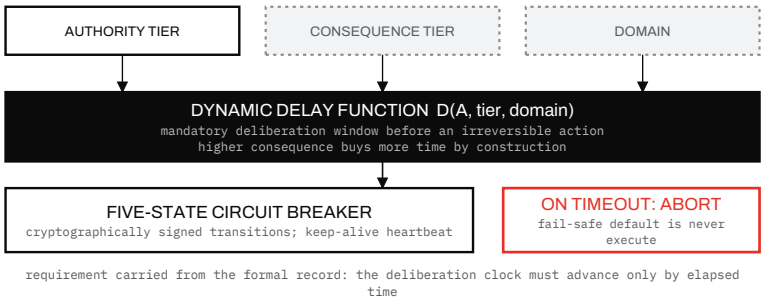

Figure 1. FLAME technical architecture. Encoding: solid = implemented; dashed = specified; gray = external; filled = enforcement point.

[HOW IT WORKS]

1. Receive the proposed action with its consequence tier, the current HMAA authority tier, and the domain.
2. Compute the deliberation window D(A, tier, domain) from the versioned policy table.
3. Arm the circuit breaker: transition to the deliberation state with a signed record and start the keep-alive heartbeat.
4. Hold the action (DELAY) while the window runs; surface it to the operator where policy requires confirmation.
5. On confirmation inside the window, transition to permit with a signed record; on timeout or lost heartbeat, transition to abort.
6. Where policy pre-authorizes autonomous execution after a bounded hold, permit only if that authority exists in the frozen envelope.
7. Record every transition, the window, and elapsed time to the ledger.

[TECHNICAL DIFFERENTIATION]

Human-on-the-loop supervision relies on an operator noticing in time; FLAME makes the time budget explicit and consequence-scaled, and it makes DELAY an outcome of the authority computation rather than an accident of workload. Static timers fail in both directions; a consequence-scaled window with abort as the default fails toward safety. The comparison is structural; no field measurement of any timer-based system is claimed.

[ASSURANCE MECHANISMS: DESIGN PROPERTY VERSUS ESTABLISHED PROPERTY]

Mechanism	Status
Timeout defaults to abort, never execute	Design property; exercised in simulation
Window scales with consequence tier	Design property; policy table versioned
Signed circuit-breaker transitions	Implemented in the simulation engine; cryptographic substrate is software, not a hardware root
Deliberation clock advances only by elapsed time	REQUIREMENT stated from a counterexample; not established for the shipped engine
No authority created by timeout	Design rule (policy-conditioned hold); exercised in AUTHREX-ABORT and ENGAGE consoles

[TECHNICAL EVIDENCE]

Evidence	Result	Source / artifact
Seeded browser simulation	Deterministic replay; consequence-scaled windows on synthetic actions	burakoktenli.com/flame-simulation
Strong-form counterexample	Outcome EXECUTE with the deliberation clock satisfied and zero real time elapsed	authrex.systems/formal-coverage
Surrogate countermodel	FLAME.tla: finding about the mechanism class, not the implemented code	formal-coverage record
Zenodo deposit	Specification v5.11, source, simulation data	DOI 10.5281/zenodo.19015618

[EVIDENCE SNAPSHOT]

EVIDENCE SNAPSHOT	
Implemented	Yes (simulation engine)
Executable artifact	Yes
Formal model	Partial (surrogate)
Simulation evidence	Yes
Hardware validation	No
Field validation	No
Independent validation	No
Government evaluation	No
Operational deployment	No

[LIMITATIONS]

- The elapsed-time requirement on the deliberation clock is stated but not established for the shipped engine; a bookkeeping clock is a published failure mode.
- Timing is idealized (about 100 Hz loop); jitter, communication delay, and processing bottlenecks on real hardware are not modeled.
- No human-factors measurement exists for whether the computed windows are usable by an operator.
- Simulation and formal-model evidence only; all data synthetic; no hardware-in-the-loop test, no field validation, no independent replication, no deployment.
- Not independently reproduced or technically verified by another party; second-environment reruns are confirmation by the same producer.

[TECHNOLOGY READINESS (SELF-ASSESSED; NO FORMAL TRA)]

CURRENT ASSESSED TRL	TRL 3 (self-assessed against the DoD nine-level scale; no formal technology readiness assessment has been performed).
BASIS	Deliberation mechanism implemented in simulation with a specified breaker and signed transitions; timing has not been characterized on any real control loop or with human operators.
EVIDENCE SUPPORTING	concept and delay function specified; simulation implementation; surrogate formal analysis
EVIDENCE REQUIRED FOR NEXT TRL	(1) Implement and check the elapsed-time clock requirement in a faithful model. (2) Software-in-the-loop timing characterization with a representative controller (Phase 4). (3) Operator-in-the-loop trial of window usability against a government decision timeline.

[INTEGRATION PROFILE]

INPUTS	Proposed action with consequence tier; HMAA tier; domain; mission clock; operator confirmation channel.
OUTPUTS	DELAY, permit, or abort; signed breaker transitions; ledger entry.
INTERFACES	Between the authority tier and execution; operator channel; ledger.
DEPLOYMENT	Software module; browser console for evaluation.
SECURITY BOUNDARY	Inside: window computation, breaker, heartbeat. Outside: clock source, operator channel. A manipulable clock defeats the window.
DATA REQUIREMENTS	None beyond a government decision timeline for Phase 2 review.

[TRANSITION ROADMAP]

Phase	Activity for this capability	Who can execute it	Exit evidence
1 Technical review	Review of the artifacts listed in the evidence table	Program alone, with a government reviewer	Reviewer findings; correction log entries
2 Government-defined simulation	Sponsor scenarios on the seeded console	Program with sponsor scenarios	Deterministic, replayable results; hash-chained records
3 Degradation and failure injection	Evidence lapse, spoofing, delayed human response	Program with sponsor parameters	False-restriction, miss, and contraction-latency figures
4 Government testbed / SIL	Integration with a representative controller or interface	Requires sponsor interface and testbed	SIL integration report
5 Hardware-in-the-loop	Representative interference; timing on real hardware	Requires sponsor laboratory or rig	Measured latencies; TRL 4 to 5 evidence
6 Relevant-environment demonstration	Independent evaluation and red team	Requires independent evaluator	Independent report
7 Program integration	Only after earlier gates succeed	Requires transition partner	Not claimed to be available

[GOVERNMENT ENGAGEMENT REQUEST (SPECIFIC)]

- Latency-budget review against a government decision timeline for a named mission (Phase 2).
- Access to a software-in-the-loop controller to measure end-to-end timing (Phase 4).
- Human-factors evaluation support to test window usability with operators.

[REVIEWER QUESTIONS THIS PROFILE ANSWERS]

What is it: see capability description. Why does DoD care: see operational relevance. Where would it fit: see integration profile. Does it exist: see evidence snapshot. What evidence supports it: see technical evidence. Has anyone independent tested it: no. Has it been tested in the real world: no. What does it rely on: see limitations and security boundary. What can cause it to fail and what does it do then: see assurance mechanisms and the published counterexamples. How mature: see the TRL block. What would the government need to provide: see the request above. Next measurable milestone: first item under evidence required for next TRL.