[CAPABILITY PROFILE · AUTHREX-07]

ERAM · ESCALATION RISK ASSESSMENT MODEL

Quantify decision-time compression and cross-domain escalation cascades in AI-enabled command and control so that escalation risk constrains authority across the pipeline rather than only being reported.

MODELED **TESTED**

Name variants: No expansion variants. ERAM is a working paper (SSRN 6176802), not a peer-reviewed publication.

PRIMARY FUNCTION	Escalation-risk quantification and decision-time-compression modeling for AI-enabled C2; feeds the consequence tier consumed by FLAME and HMAA; implemented as the ERAM Strategic Command simulation (6 scenarios x 100 = 600 Monte Carlo runs) with Merkle provenance chains and a sealed decision ledger.
MISSION AREAS	AI-enabled command-and-control experimentation; strategic-level escalation studies; wargame support; joint all-domain C2 modeling.
EVIDENCE ANCHORS	SSRN working paper 6176802; ERAM Strategic Command console (SIM-01); formal-coverage record: ERAM.tla surrogate countermodel with a published strong-form counterexample (SilentDrop).
CURRENT STANDING	Working paper and seeded simulation; model-based with documented parameter assumptions; not validated against operational data; not peer reviewed; not independently reproduced.
SELF-ASSESSED TRL	TRL 2 to 3 (analytical model with a simulation demonstrator). No formal TRL determination has been performed.
INTELLECTUAL PROPERTY	U.S. provisional application 64/110,225, filed July 13, 2026 (receipt 78660113), reported as filed; unexamined; confers no enforceable rights.

[MISSION PROBLEM]

Local autonomous actions can cascade across domain boundaries and compress decision time at the strategic level: a defensive action in one domain becomes a cue in another, and each automated step removes time from the next human decision. Governance layers that look only at the local action miss the cascade. The consequence is escalation by diffusion, where no single decision was unauthorized but the aggregate outran every deliberate decision cycle.

[OPERATIONAL RELEVANCE]

- AI-enabled command and control where recommendations must carry a consequence tier that reflects cross-domain coupling.
- Strategic-level escalation studies and wargame support with reproducible, seeded runs.
- Joint all-domain C2 experimentation environments (APEX JADC2 Theater and MDO Digital Twin simulations exercise the coupling).
- Audit reconstruction: the sealed ledger conventions used across the program originate here.

[CAPABILITY DESCRIPTION]

ERAM is the cross-cutting monitor of the pipeline. It consumes pipeline state (tiers, windows, recovery phases) and mission context, and models how local actions propagate across domains and compress decision time. Its output is a consequence tier that FLAME uses to size the deliberation window and HMAA uses in tier computation, so escalation risk constrains authority instead of only reporting it. The ERAM Strategic Command simulation runs six scenarios with 100 Monte Carlo runs each under a seeded generator, with formal properties checked on the model and a Merkle provenance chain over the run record. The formal record adds a requirement: ledger integrity needs an independent decision count to reconcile against, because a valid hash chain proves only that what was recorded was not altered, not that nothing was silently dropped.

[ERAM · ARCHITECTURE, ASSURANCE, EVIDENCE]

[TECHNICAL ARCHITECTURE]

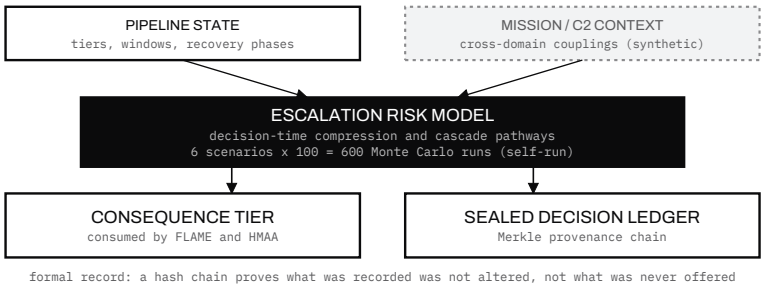

Figure 1. ERAM technical architecture. Encoding: solid = implemented; dashed = specified; gray = external; filled = enforcement point.

[HOW IT WORKS]

1. Ingest pipeline state from every stage and the mission or C2 context (synthetic in the demonstrator).
2. Model cross-domain couplings and compute decision-time compression along each escalation pathway.
3. Derive the consequence tier for the pending action class and publish it to FLAME and HMAA.
4. Run the seeded Monte Carlo campaign for the scenario under evaluation and compute the escalation-risk distribution.
5. Seal the run record with a Merkle chain and reconcile the ledger against an independent decision count.
6. Report the distribution and the constraint applied; the outcome itself is issued by HMAA.

[TECHNICAL DIFFERENTIATION]

Escalation-risk analysis is usually a study product delivered after the fact; ERAM places it inside the runtime pipeline as a constraint on authority. Against anomaly detection it converts a signal into a consequence tier that changes how much time and authority the next action receives. The model is a working paper and the differentiation is architectural; no validation against operational escalation data exists.

[ASSURANCE MECHANISMS: DESIGN PROPERTY VERSUS ESTABLISHED PROPERTY]

Mechanism	Status
Consequence tier constrains FLAME and HMAA	Design property; implemented in the integrated consoles
Seeded, reproducible Monte Carlo campaign	Implemented; bit-exact within the packaged build; not independently reproduced
Merkle provenance over the run record	Implemented in the console
Ledger completeness (nothing silently dropped)	NOT established: SilentDrop counterexample; requirement for an independent decision count recorded
Formal properties on the model	Checked in the ERAM Strategic Command console at stated bounds; ERAM.tla is a surrogate

[TECHNICAL EVIDENCE]

Evidence	Result	Source / artifact
Working paper	SSRN 6176802 (working paper, not peer reviewed)	SSRN
ERAM Strategic Command simulation	6 scenarios; 600 Monte Carlo runs; formal properties; Merkle chain; seeded PRNG	authrex.systems (SIM-01)
Strong-form counterexample	SilentDrop fires, ledger is empty, shadow is non-empty, hash chain valid	authrex.systems/formal-coverage
Earlier wording "validated across all 600 Monte Carlo runs"	Removed from the live site; results are demonstrated in simulation, not validated	Correction log CL-05

[EVIDENCE SNAPSHOT]

EVIDENCE SNAPSHOT	
Implemented	Yes (console)
Executable artifact	Yes
Formal model	Partial (surrogate)
Simulation evidence	Yes
Hardware validation	N/A
Field validation	No
Independent validation	No
Government evaluation	No
Operational deployment	No

[LIMITATIONS]

- Model-based with documented parameter assumptions; not validated against operational or historical escalation data.
- Working paper status: not peer reviewed.
- Ledger completeness cannot be established from the chain alone; the independent decision count is a requirement, not a shipped mechanism.
- Six scenarios are author-defined; no government scenario has been run.
- 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 2 to 3 (self-assessed against the DoD nine-level scale; no formal technology readiness assessment has been performed).
BASIS	Analytical model with a reproducible simulation demonstrator; the critical function has not been validated against any real data, which limits the assessment to concept and analytical proof of concept.
EVIDENCE SUPPORTING	concept formulated (working paper); simulation demonstrator
EVIDENCE REQUIRED FOR NEXT TRL	(1) Peer-reviewed submission with sensitivity analysis. (2) Government-defined C2 scenarios on the console (Phase 2). (3) Comparison against a sponsor-provided escalation case study or wargame record.

[INTEGRATION PROFILE]

INPUTS	Pipeline state; mission and C2 context; scenario definitions; seed.
OUTPUTS	Consequence tier; escalation-risk distribution; sealed run record.
INTERFACES	FLAME and HMAA inputs; ledger; console export.
DEPLOYMENT	Standalone browser console (client-side, no network); analytical tool rather than an operational decision aid.
SECURITY BOUNDARY	Inside: model, chain. Outside: context sources. The chain does not prove completeness.
DATA REQUIREMENTS	Synthetic; wargame or exercise records under sponsor control would enable validation.

[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)]

- Working-paper review by a C2 or escalation-dynamics research organization.
- A government wargame or exercise dataset (releasable) for model validation.
- Integration into a sponsor's C2 experimentation environment as an analytical monitor, not a decision aid.

[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.