

[AUTHORITY LIFECYCLE GOVERNANCE]

AUTHREX

EXECUTIVE CATALOG

Runtime authority, assurance, and verification for AI-enabled and autonomous systems. Seven governance frameworks, one executable governance kernel, two frozen research consoles, six software application profiles, twelve hardware reference designs and fourteen catalogued simulation records, documented at their actual standing with evidence, limitations, and a gated government evaluation path for each.

READ THE CORRECTIONS FIRST → PAGE 5

TRL REGISTER → PAGE 7

ENGAGEMENT MATRIX → PAGE 10

7

GOVERNANCE
FRAMEWORKS

1

EXECUTABLE
KERNEL

12

HARDWARE DESIGNS
(NONE BUILT)

14

CATALOGUED
SIM RECORDS

3

PEER-REVIEWED
ARTICLES

8

PROVISIONALS
REPORTED AS FILED

0

INDEPENDENT
VALIDATIONS

[01] · WHAT EXISTS

Specifications, formal models checked at stated bounds, an executable governance kernel with a container test suite, deterministic seeded consoles with hash-chained records, bill-of-materials-level hardware designs, three peer-reviewed journal articles, and eight provisional applications reported as filed.

[02] · WHAT DOES NOT

No hardware built. No hardware-in-the-loop test. No field validation. No independent replication. No deployment, certification, accreditation, or government adoption. One favorable AFRL white-paper evaluation on record; not funded; not an endorsement.

[03] · THE ASK

Phase 1 technical review of models, code, and evidence (no hardware, no funding); Phase 2 government-defined scenarios on the seeded consoles; Phase 3 degradation and failure injection; then a sponsor testbed and hardware-in-the-loop.

INDEPENDENT RESEARCH PORTFOLIO. NOT A GOVERNMENT PUBLICATION; NOT AFFILIATED WITH, ENDORSED BY, OR PRODUCED FOR ANY GOVERNMENT AGENCY. NO CLAIM OF ADOPTION, CONTRACT, CERTIFICATION, ACCREDITATION, OR CLEARANCE IS MADE. EVERY RESULT IS SELF-REPORTED AND NOT INDEPENDENTLY VERIFIED. DISTRIBUTION: PUBLIC RELEASE (OPEN INFORMATION). LICENSE: CC BY 4.0. NO FORMAL DOD DISTRIBUTION STATEMENT IS ASSIGNED.

BURAK OKTENLI, MBA · PRINCIPAL INVESTIGATOR · INDEPENDENT RESEARCHER, AUTHREX RESEARCH PROGRAM, WASHINGTON, DC · ORCID 0009-0001-8573-1667
AUTHREX.SYSTEMS · BURAKOKTENLI.COM · INFO@BURAKOKTENLI.COM · VERSION 6.1 · SEPTEMBER 4, 2026 · SUPERSEDES AUTHREX-CAT-2026-002 (V5, 1 SEP 2026)

AUTHREX AT A GLANCE

This page stands alone if forwarded without the catalog. A reader with thirty seconds should find who, what, why, proof, maturity, and the ask.

7	1	12	14	3	8	0
FRAMEWORKS	EXECUTABLE KERNEL	REFERENCE DESIGNS, 0 BUILT	CATALOGUED SIM RECORDS	PEER-REVIEWED ARTICLES	PROVISIONALS FILED	INDEPENDENT VALIDATIONS
WHO	A single principal investigator (Burak Oktenli, MBA; B.S. Computer Science, University of South Florida, 2020; M.B.A., Lynn University, 2026; M.P.S. Applied Intelligence, Georgetown University, in progress; prior industry experience in production data engineering). No standing engineering team, hardware laboratory, or independent test authority inside the program.					
WHAT	A runtime authority and assurance architecture that makes machine permissions explicit, enforceable, contractible, and auditable as operational evidence changes. Each of position, sensing, communications, and software integrity carries its own validity state; the system may execute only what every required condition admits; a strong source never compensates for a failed one; on failure the admissible set narrows at once and widens only by authorization. Every outcome (EXECUTE, DELAY, HANDOFF, ABORT) is written to a hash-chained record.					
WHY	The gap appears where two conditions hold together: a machine can select and execute an action faster than its supervisor can evaluate it, and the sensing that justifies the action can be degraded, jammed, or deceived. Uncrewed air systems under GNSS denial, counter-UAS engagement decisions, ground robotics, spacecraft under light delay, and agentic AI on networks all sit inside that gap.					
PROOF	HMAA authority automaton model checked in TLA+ (23,748 distinct states at depth 9; 5 invariants and 1 liveness held; 2 vacuous at the bound; discrete automaton only). AUTHREX-ABORT: two TLA+ modules across 21 and 35 configurations, 105 tests, 1,000,000-run soak. Governance Kernel: 83 tests passed in a container. Three peer-reviewed journal articles (Springer Nature x2, IJCSS); one favorable AFRL white-paper evaluation on record (CANVAS, June 2026; not funded; not an endorsement). All self-run on synthetic data.					
MATURITY	Research stage. Self-assessed TRL 2 to 4 by artifact; no formal TRL determination. No hardware-in-the-loop test, no field validation, no independent replication, no deployment, no certification, no government adoption. Twelve hardware designs exist at bill-of-materials level; none is built.					
ASK	Phase 1 technical review of models, code, and evidence (no hardware, no funding); Phase 2 government-defined scenarios on the seeded consoles; Phase 3 degradation and failure injection; then a government testbed and hardware-in-the-loop. Specific needs per capability: recorded sensor streams, a SIL flight or vehicle stack, a platform safety-interface specification, an independent evaluator, and a sponsor-stated mission use case.					
CONTACT	Burak Oktenli, MBA · AUTHREX Research Program · info@burakoktenli.com · authrex.systems · burakoktenli.com · ORCID 0009-0001-8573-1667					

Self-reported. Not independently verified. Not a government publication. No adoption, contract, certification, or endorsement is claimed. The evidence register and claims register at the back of the catalog are the record of what exists.

EXECUTIVE OVERVIEW

What the program does, what it has built, and what it is asking for.

[WHAT THE PROGRAM DOES]

The AUTHREX Research Program develops runtime authority governance and assurance for autonomous and AI-enabled systems. It sits beside the mission autonomy rather than inside it and separates what a system can technically do from what it remains authorized to do at this moment, on this evidence. The program has produced seven governance frameworks that together form a seven-stage authority lifecycle, one executable governance kernel for agentic AI, two frozen research consoles (AUTHREX-ABORT for evidence-conditioned authority lapse and AUTHREX-DA for data admissibility), six software application profiles, twelve hardware reference designs, and a seeded, deterministic simulation portfolio, all supported by a public disclosure corpus with permanent identifiers.

[MAJOR TECHNICAL DOMAINS]

Autonomous systems (air, ground, maritime, space); AI assurance and agentic AI; mission assurance and command and control; cyber-physical systems and operational technology; formal methods; resilient PNT as an evidence problem; hardware roots of trust (reference designs only).

[CENTRAL ENGINEERING POSITION]

Authority is granted in advance by a human, bounded by an explicit envelope, and conditioned on evidence. When the evidence degrades, the admissible set of actions narrows at once; when it fails, authority lapses and irreversible actions are latched down. Authority widens only by an authorized transition, never by timeout and never by recovery of a signal alone. A confident model cannot compensate for invalid navigation.

[PORTFOLIO MATURITY, STATED PLAINLY]

- Frameworks: specified, simulated on synthetic data, and formally modeled at stated bounds where stated; HMAA_Refine is faithful on the discrete side; two of twenty TLA+ modules are faithful to implemented code. Self-assessed TRL 3.
- Governance Kernel: implemented, container-tested (83 passed, 1 skipped), not released, not independently validated. Self-assessed TRL 3 to 4.
- AUTHREX-ABORT: two TLA+ modules model checked at stated bounds, 105 implementation tests, open gates published. Self-assessed TRL 3.
- Hardware: twelve bill-of-materials-level reference designs; none built; no physical measurement. Self-assessed TRL 2, with BLADE-AV, AGENT-HSM and SWARM at 2 to 3 on simulation or emulator evidence.
- Independent validation: none. Government evaluation: one favorable AFRL white-paper evaluation memorandum (June 2026), not funded, not an endorsement. Deployment: none.

[GOVERNMENT ENGAGEMENT SOUGHT]

A technical review of the models, code, and evidence; a government-defined mission use case with an explicit confirm-or-abort tier; government-defined scenarios run on the seeded consoles; degradation and failure injection; then access to a software-in-the-loop stack, a hardware-in-the-loop environment, and an independent evaluator. Candidate vehicles that fit the maturity: research agreement or CRADA, BAA or SBIR/STTR topic response, testbed collaboration with an FFRDC, UARC, or laboratory, subcontract to a platform integrator, or a Phase 1 to 3 pilot. No mechanism is stated to be available or guaranteed.

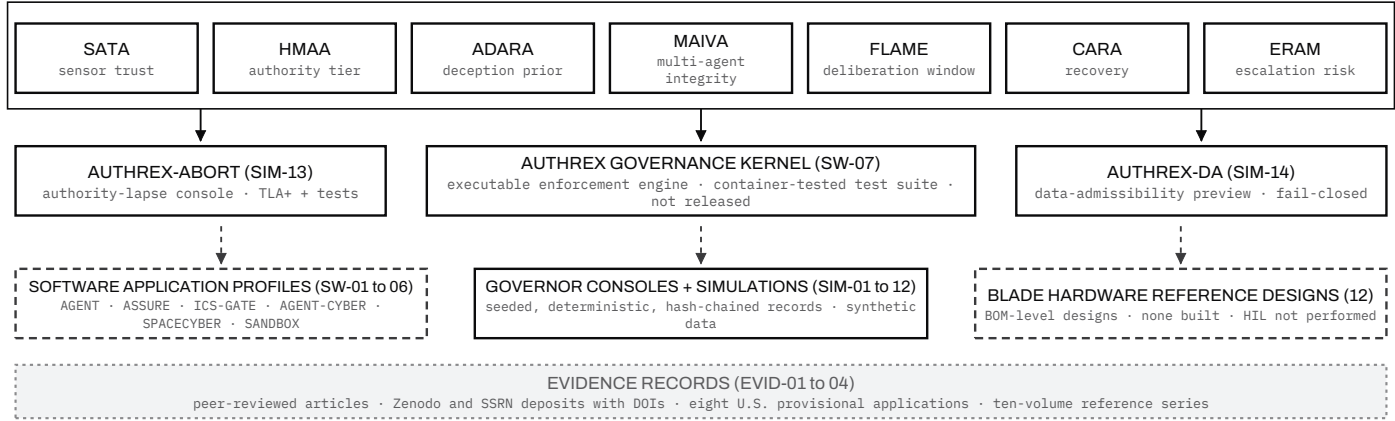
[WHAT THIS CATALOG DOES NOT CLAIM]

- That formal verification of a model proves any deployed implementation; bounded model checking of a configured finite model is not a proof of the physical system.
- That simulation equals field validation; every result here is simulation on synthetic data unless hardware is named, and no hardware is built.
- That any government organization has adopted, contracted, evaluated (beyond the AFRL record), certified, or endorsed the work.
- That agreement across evidence sources is truth; correlated false information is an open limitation (NEG-3 in AUTHREX-DA; identical-lie counterexample in ADARA).

PORTFOLIO AT A GLANCE

Forty-four register items in four families. What each is, what it addresses, how mature, what the evidence is, how it would connect, and what the government would need to provide.

GOVERNANCE FRAMEWORKS (7) · specified, simulated, formally modeled where stated



Encoding: solid = executable artifact exists; dashed = specification or reference design; gray = documentary record; arrows = the frameworks are instantiated by the artifact below.

Figure 1. Portfolio structure. The frameworks define the authority logic; the kernel, consoles, and simulations execute it in software; the hardware family specifies it for platforms; the evidence records document it.

Capability	Mission problem	Primary function	Current maturity	Evidence	Integration path	Government need
SATA	Sensor integrity cannot be assumed under jamming, spoofing, compromise	Per-source trust scalar from four diagnostics with attestation	TRL 3 (self); published protocol	JHSS 2026 article; Zenodo; simulation	Sensor-ingest boundary module	Recorded sensor streams; TPM node
HMAA	Binary permit/deny cannot express partial trust or degrade gracefully	Graded A3 to A0 tier with asymmetric hysteresis	TRL 3 (self); TLA+ at bounds	23,748 states depth 9; 98 tests	Enforcement point on the actuator path	Scenario with confirm-or-abort tier; SIL stack
ADARA	Manipulated perception inherited by governance	Deception prior modulating authority downward	TRL 3 (self)	Simulation; one faithful module; counterexample	Parallel tap on SATA output	Deception scenarios with ground truth; red team
MAIVA	One compromised agent or contaminated majority carries the outcome	Byzantine-resilient trimmed weighted median; CUSUM	TRL 3 (self); open defect	Faithful quorum model; 37 self-tests	Aggregator before the tier	Tolerable-compromise assumption; multi-vehicle SIL
FLAME	Veto not exercisable before an irreversible outcome	Consequence-scaled deliberation window; abort on timeout	TRL 3 (self)	Simulation; counterexample; requirement	Between tier and execution	Decision timeline; operators for a trial
CARA	Recovery short-circuited by one favorable signal	GREP recovery with non-compensatory terminal gate	TRL 3 (self)	Self-run enumeration; simulation	Invoked on ABORT or lockout	Recovery policy for a platform class
ERAM	Local actions cascade and compress strategic decision time	Escalation-risk model feeding the consequence tier	TRL 2 to 3 (self); working paper	SSRN 6176802; 600-run console	Cross-cutting monitor	Releasable exercise record
Governance Kernel	Agents hold whatever privilege their keys hold	Tier gating, approval queue, signed ledger at the tool boundary	TRL 3 to 4 (self); not released	83 tests passed in container	Proxy / runtime wrapper (MCP)	External reviewer team; tool inventory and policy
AUTHREX-ABORT	Fallbacks add capability when the system is least trustworthy	Evidence-conditioned authority lapse with latch	TRL 3 (self); gates open	2 TLA+ modules; 105 tests; soak	Abstract platform adapter	Safety-interface specification; T&E; adoption
AUTHREX-DA	Poisoned or repaired data admitted without an authority decision	Fail-closed admissibility verdicts with reason codes	TRL 2 to 3 (self); WP-0 blocked	101 tests; single-seed root; NEG-1 to 5	Pipeline stage (design)	Labeled poisoned records
Software profiles (6)	Same gap in agent, OT, cyber, space, and T&E; settings	Pipeline applied per setting; governor consoles	TRL 2 to 3 (self)	Specifications; consoles	Runtime wrapper or middleware	Select one profile for Phase 2
BLADE designs (12)	Hardware enforcement and roots of trust per domain	BOM-level governance nodes	TRL 2 (self); none built	Zenodo deposits; simulation; HSM emulator	Mission-system module (design)	Bench, HIL rig, range; integration partner

CORRECTIONS FIRST: WHAT CHANGED SINCE VERSION 5

The program applies withdrawals in place and records them beside the original. A reader who holds an earlier document should read this page before any capability page.

ID	Item	Earlier statement	Statement in this edition
CL-01	HMAA model-checking figure	Earlier figure of 48,751 states appears in the June 2026 two-page summary and older pages	Withdrawn; canonical figure is 23,748 distinct states at depth 9 with its scope clause, never stated without it
CL-02	SATA archived TLC run	formal-coverage.html still states a separate archived result of 523 distinct states at depth 9; the specification, configuration and log for that run could not be produced on request	Withdrawn from this catalog pending recovery of the artifacts; the as-published article figure (18,892 states) is cited only as published
CL-03	Rover 350-run and UAV 250-run counts	Withdrawn 7 Aug 2026 as unreproducible; still present on burakoktenli.com pages (evaluation.html, platform pages) and in the July white paper and two BAA white papers	Replaced by "seven fault scenarios" and "five adversarial scenarios" demonstrated in simulation; the aggregate "2,800+ runs" is not carried
CL-04	"Physically assembled" testbeds	July 2026 white paper describes the rover and UAV testbeds as physically assembled; the register records no platform as built	This edition states not built, physical build not yet documented
CL-05	Validation language on live pages	"Simulation Validated", "Validated in rover testbed", "Six Validated Scenarios", "ALL VALIDATIONS PASSED" appeared on framework and platform pages	Removed from the 1 Sep 2026 deploy; this edition uses demonstrated in simulation
CL-06	Simulation count: 24 versus 14 versus 37	V5 says twenty-four seeded simulations (counting governor consoles); the register has 14 SIM rows; the sites state 37 browser simulations (19 plus 18)	This edition uses 14 catalogued records (SIM-01 to 14) and states the other counts with their definitions; never interchangeable
CL-07	Name expansions	HMAA (Allocation, Architecture, Hierarchical Mission Authority Architecture); SATA (Attestation and Trust Anchoring); ADARA (Risk Architecture); MAIVA (Verification); FLAME (Flash War Latency Architecture); CARA (Control Authority Regulation; Cognitive Authority Recovery in the Zenodo title)	Canonical program names used; variants listed on each capability page; original artifact titles preserved
CL-08	Provisional receipts	V5 lists the three July 2026 filings as receipt pending; the patents page (verified 1 Sep 2026) carries receipts 78658472, 78659350, 78660113	Receipts carried in this edition
CL-09	Peer-reviewed article count	V5 states one peer-reviewed article; two further articles published 24 Aug and 31 Aug 2026	Three stated, with the caveat that one is not an AUTHREX technical paper
CL-10	BLADE-UUV and domain counts	The July white paper and the authrex.systems index list BLADE-UUV; the register has no such entry; domain counts vary (nine, ten, eleven) across properties	BLADE-UUV not carried (no deposit, no register row); the catalog counts twelve register platforms and does not state a domain count
CL-11	Standards-mapping overclaims	standards-mapping.html states "no unsafe state reachable from any initial state under any sensor input" and describes AUTHREX as "required by" MIL-HDBK-516C sections	Not carried; this edition uses mapped by the author, design target, architecturally relevant
CL-12	Superseded governance formulations	June 2026 wording; scores the evidence; grants authority in proportion; narrows in proportion; so the human stays in command; "narrows, and does not widen on its own"	Replaced by the per-source evidence-vector, non-compensatory-gate, admissible-set formulation; "narrows at once; widens only by authorization"
CL-13	Absolute hold rule	Earlier form: every irreversible action requires confirmation and cannot execute without it	Policy-conditioned form: where policy requires confirmation, execution occurs only after it; if absent at hold expiry the predeclared outcome is delay, handoff, or abort; pre-authorized autonomous execution must already exist in the envelope; timeout never creates authority

The full correction log, including items pending owner decision and pending page fixes on the public sites, is at the back of the catalog.

WHERE THE PROBLEM APPEARS

The problem is not autonomy in general. It appears where a machine can act faster than its supervisor can evaluate and the sensing that justifies the action can be degraded, jammed, or deceived. Either condition alone is covered by existing practice. Both together is the gap.

Where both conditions hold	Decision window	Degradable evidence	Relevant capabilities	What the portfolio offers today
Uncrewed air systems in GNSS-denied or contested electromagnetic conditions	Seconds	GNSS, RF, datalink	SATA, HMAA, ABORT	Trust falls; tier contracts at once; navigation failure is a hard veto; latch on lapse (simulation)
Counter-UAS engagement decisions	Seconds	Radar track, RF identification, EO/IR	ADARA, FLAME, HMAA, BLADE-CUAS	Deception prior extends the deliberation window; HANDOFF at A1 (simulation)
Ground robotic movement and convoy operations	Seconds to minutes	Lidar, vision, positioning	SATA, HMAA, CARA, Rover design	Speed and route constrained at A2 and A1; A0 (revoked) safe-stop; GREP recovery (simulation)
Spacecraft operations under light delay	Supervision impossible in real time	Star tracker, ranging, telemetry	FLAME, CARA, SPACECYBER, BLADE-SPACE	DELAY then ABORT to A0 if evidence is not refreshed; staged recovery (design and simulation)
Agentic AI acting on networks and infrastructure	Milliseconds	Tool output, retrieved content, provenance	Governance Kernel, AUTHREX-AGENT, DA	Tier enforcement and approval gating at the tool boundary; signed evidence (container-tested)
Operational technology and critical infrastructure	Seconds to minutes	Sensor plausibility, command provenance	ICS-GATE, BLADE-INFRA-OT, FLAME	Mandatory deliberation before control writes; fail-closed boundary adjudication (design and console)
Multi-platform and swarm operations	Seconds	Inter-agent attestations, RF	MAIVA, BLADE-SWARM	Trimmed weighted median; safe-halt by default under denied RF (simulation; open quorum defect)

OPERATIONAL PROBLEM TO CAPABILITY

Operational problem	Capability	How it helps (as demonstrated in simulation)
Sensor confidence degrades	SATA, HMAA	Trust scalar falls; authority tier contracts immediately; upgrade is delayed by hysteresis and never restored by timeout
Communications lost	FLAME, CARA, ABORT	Latency budget holds the decision (DELAY); loss of required evidence contracts the admissible set; pre-authorized contingencies only
Target evidence becomes stale	FLAME, HMAA	Deliberation window and evidence age gate the irreversible action; HANDOFF at A1 or ABORT to A0
Agent exceeds assigned permission	Governance Kernel, AUTHREX-AGENT	Tier enforcement and human approval gating at the agent runtime; signed audit evidence
Platform enters degraded mode	CARA	Deterministic GREP recovery with a non-compensatory terminal gate
Audit reconstruction required	All consoles	Hash-chained record with deterministic replay; a reviewer re-runs a result from the same packaged build
Data of uncertain provenance offered	AUTHREX-DA	ADMIT, ADMIT_FLAG, QUARANTINE, or REJECT with reason code; never repaired
Safety invariant must hold	HMAA, ABORT (TLA+)	Invariants held at stated bounds; bounded model checking is not a proof of the physical system

The applications above are potential mission applications. No adoption, contract, or evaluation by any government organization is claimed, and no specific DoD program is stated to be using or considering the work.

WHERE THE AUTHORITY LAYER SITS

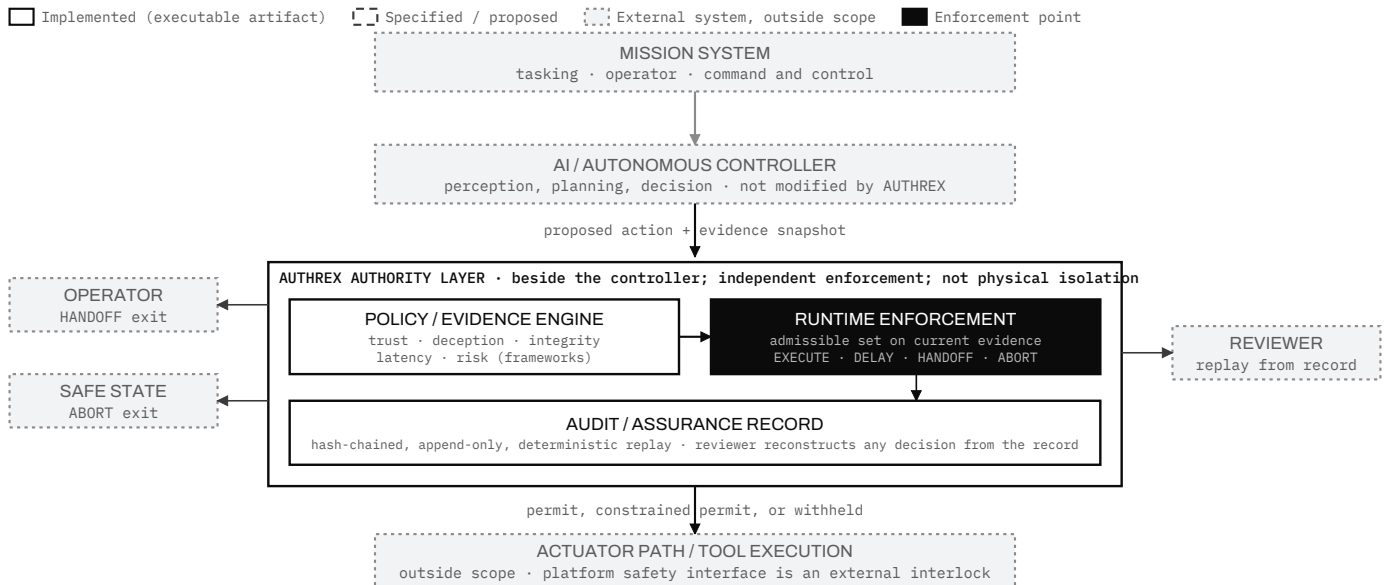



Figure 2. Layered view and trust boundary.



HMAA is the only component that issues an authority tier; every outcome (EXECUTE, DELAY, HANDOFF, ABORT) is issued from that one place.

Figure 3. Seven-stage lifecycle.

TECHNOLOGY READINESS REGISTER

Self-assessed against the DoD nine-level TRL scale (DoD Technology Readiness Assessment guidance; software TRL descriptions). No formal technology readiness assessment has been performed by any authority. A simulation is not TRL 5 to 6; a formal model is not an operational prototype; a reference design is not a breadboard.




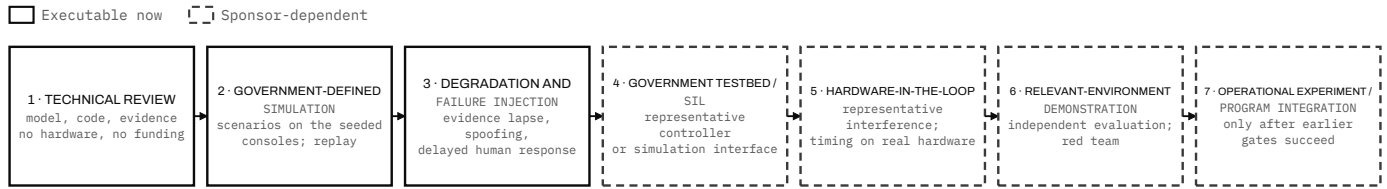
Figure 4. Self-assessed readiness by family. Near-term maturity gates differ by artifact and are listed in the register below. Hardware-in-the-loop under representative interference is a later transition gate for the platform-facing capabilities, not the next gate for every family.

Item	TRL (self)	Rationale	Missing evidence	Next milestone
SATA	3 (article states 4)	Published protocol; reference implementation; seeded simulation; formal check as published. The article self-assesses TRL 4; this register holds TRL 3 pending laboratory component validation	Laboratory component validation with attestation hardware; independent replication	HIL on ROS 2 with injected degradation
HMAA	3	TLA+ at stated bounds; 98-test package; simulation	SIL with a representative controller; closure of two vacuous properties; Obligation 6	SIL integration with a flight or vehicle stack
ADARA	3	Simulation; one faithful formal module; published counterexample	Repair shipped and checked; scenarios with ground truth	Government-defined deception scenarios
MAIVA	3	Self-tested implementation; faithful quorum model located an open defect	Repair shipped; multi-agent SIL	Repaired quorum re-check
FLAME	3	Simulation; surrogate formal analysis; stated clock requirement	Timing on a real loop; operator trial	Elapsed-time clock requirement checked in a faithful model
CARA	3	Simulation; self-run enumeration	Independent reproduction; nonce-bound token shipped	Recovery trial on SIL platform
ERAM	2 to 3	Working paper; seeded demonstrator	Validation against any real data; peer review	Peer-reviewed submission
Governance Kernel (SW-07)	3 to 4	Executable prototype with container test suite	External reproduction; security assessment; representative workload	Structured external review
Software profiles (SW-01 to 06)	2 to 3	Specifications with governor consoles	Any implementation beyond the console	Select one profile for Phase 2
AUTHREX-ABORT (SIM-13)	3	Two TLA+ modules; 105 tests; soak; mutation campaign	Refinement gate; negative controls; hazard and cyber packages	Engine-to-InterfaceRecord conformance gate and the B14 and B2 negative controls closed
AUTHREX-DA (SIM-14)	2 to 3	Local tests; single-seed root	WP-0; DA-P1 to P9; preregistered campaign	WP-0 closure
Governor consoles (SIM-08 to 12)	3 (as tools)	Deterministic consoles with reference roots	Signed releases; independent replay	Signed release of TEAM and REDLINE
BLADE-AV	2 to 3	BOM-level design; 12 attack scenarios in simulation	Bench integration of the fail-safe relay; drive-by-wire rig	Bench relay integration
BLADE-AGENT-HSM	2 to 3	Design plus adversarial emulator (275/275)	Silicon; measured timing	Prototype board
BLADE-SWARM	2 to 3	TLA+ at stated bounds; simulation at N up to 500	Any physical node	Multi-node SIL
BLADE-EDGE, AV, MARITIME, INFRA, INFRA-OT, CUAS, FINANCE, Rover, UAV	2	BOM-level reference designs; simulation data	Any physical build or measurement	First prototype assembly (BLADE-EDGE or Rover)
BLADE-SPACE	2	Preliminary design; V&V; specified, not executed	Executed V&V; orbital simulation campaign	Execute specified V&V;

Reading rule for TRL 3 versus 4 in this register: TRL 3 requires analytical and experimental proof of concept of the critical function, which simulation on synthetic data with a formal model satisfies; TRL 4 requires validation of a component in a laboratory environment, which for software means a prototype validated beyond the author's development environment against a representative interface or workload. Only the Governance Kernel approaches that threshold, and it is stated as 3 to 4 rather than 4.

TRANSITION ROADMAP

Current state, integration prototype, relevant environment, mission experiment, program integration: seven gated phases. Phases 1 to 3 require no hardware and no deployment funding.



CAN START NOW: no hardware and no deployment funding required

REQUIRES A SPONSOR: use case, interface, testbed or range access, independent evaluator

Each phase is a gate. A phase proceeds only after the prior one succeeds and its evidence is published to the register. Solid boxes are phases the program can execute unaided today; dashed phases depend on government-furnished scenarios, interfaces, test environments, or an independent evaluator. No phase implies that any organization has agreed to sponsor it.

Figure 5. Gated evaluation path.

30 / 60 / 90 DAY PILOT STRUCTURE (PHASES 1 TO 3)

First 30 days	Days 31 to 60	Days 61 to 90
Architecture review; interface definition; government use-case selection; baseline scenario construction; integration planning	Prototype integration with a representative controller or simulation interface; government-defined tests; failure injection; metrics collection	Demonstration; independent evaluation; technical report; transition recommendation

WHAT EACH SIDE PROVIDES

GOVERNMENT, REQUIRED	Mission scenario; representative controller or simulation interface; evaluation metrics.	AUTHREX PROVIDES	Software and reference architectures; TLA+ models and enumeration artifacts; seeded consoles and test harness; assurance logic and integration support; evidence reports; source access under invited review.
GOVERNMENT, OPTIONAL	Test data; operational constraints; access to a testbed or range; an independent evaluator.	NOT OFFERED	Deployed hardware, certified software, or operational support, none of which exist.

SUCCESS METRICS A SPONSOR COULD ADOPT

Unauthorized transition prevention; authority-contraction latency; deterministic state behaviour under replay; failure-response correctness; false-restriction rate (authority reduced in a benign environment); false-negative rate; recovery behaviour; audit reconstruction success; computational overhead and share of the latency budget consumed; integration burden; operator comprehension of authority state.

TRANSITION PATHWAYS THAT FIT THE CURRENT MATURITY

Research agreement or CRADA; university research collaboration; BAA or SBIR/STTR topic response; testbed collaboration with an FFRDC, UARC, or laboratory; subcontract or prime integration for a specific platform; government pilot at Phases 1 to 3; consortium membership through an OTA. No mechanism is stated to be available or guaranteed; each depends on a sponsor identifying a use case. Eligibility for any vehicle is not asserted.

ACQUISITION-READY INFORMATION

Field	Value	Field	Value
Legal entity	Formation in progress; no entity identifiers available	UEI / CAGE	Not yet issued
SAM status	Not registered	NAICS	Not assigned
Small-business status	Not determined	Contract vehicles	None
Security clearance	Not provided	Facility clearance	None
Program staffing	One principal investigator	Teaming on award	Hardware integration, HIL, red team, certification evidence, and external formal-methods review would be teamed; named as gaps, not claimed as capabilities

GOVERNMENT ENGAGEMENT MATRIX

For every capability: the kind of organization that fits, the specific contribution needed, the evaluation proposed, and the transition objective. "We are looking for DoD partners" is not a request; these are.

Capability	Desired organization type	Government contribution needed	Proposed evaluation	Transition objective
SATA	Service laboratory (AFRL, ARL, NRL); FFRDC	Recorded sensor streams with degradation episodes; ROS 2 platform with TPM node	Phase 1 to 3 then HIL	Component validated in a laboratory environment (TRL 4)
HMAA	Formal-methods-capable laboratory; T&E; organization	Mission scenario with explicit confirm-or-abort tier; SIL flight or vehicle stack	Independent review of TLA+ and tests; SIL integration	TRL 4; policy thresholds set by a sponsor
ADARA	EW or counter-UAS test organization	Deception scenarios with ground truth; red team	Phase 2 simulation; adaptive red team	Calibrated false-restriction and miss rates
MAIVA	Swarm or multi-platform program office	Tolerable-compromise assumption; multi-vehicle SIL	Adversarial agreement evaluation	Repaired quorum validated in SIL
FLAME	C2 or engagement-chain program office; human-factors laboratory	Government decision timeline; operators for a usability trial	Latency-budget review; operator-in-the-loop	Windows usable against a named timeline
CARA	Platform safety authority (space, maritime, ICS)	Recovery policy for a named platform class	Recovery-path review; SIL lockout trial	Recovery validated in SIL
ERAM	C2 research organization; wargaming center	Releasable exercise or wargame record	Model validation against the record	Validated analytical monitor
Governance Kernel	Cyber test organization; DIU or service software factory	Tool inventory and approval policy; reviewer team	External reproduction; sandbox evaluation; security assessment	Independently reproduced; assessed for release
AUTHREX-ABORT	DT&E; or OT&E; organization; UAS program office	Platform safety interface specification; hazard-analysis review	Adopt the evidence-lapse pattern as a test case	Bench demonstrator with a real safety interface
AUTHREX-DA	Data-integrity or federated-learning program	Labeled records with poisoned entries	Admission evaluation	WP-0 closed; multi-seed campaign
BLADE family	Laboratory with hardware integration capacity; prime contractor	Bench, HIL rig, range access; interface specifications	Design review then prototype assembly	First platform built and measured

SEVEN LEVELS, FROM A CONVERSATION TO A PILOT

Level	What it involves	What the government receives	What is acquired
Technical exchange	30 to 60 minute architecture discussion	Answers on scope, evidence, limitations	Nothing
Use-case review	Apply the architecture to a specific mission scenario	Written mapping of the scenario to the authority lifecycle	Nothing
Independent technical assessment	Code, formal artifacts, test evidence provided for review	Source access under invited review; TLA+ models; enumeration artifacts	Access for review
Simulation evaluation	Government-defined scenarios on the seeded consoles	Deterministic, replayable results; hash-chained records	Access for review
Pilot integration	Integrate with a representative platform or simulation interface	Phases 1 to 3; technical report	Research and prototype integration engineering, not a product
Research collaboration	Advance formal methods, assurance, or mission applications	Joint artifacts; shared replication packages	Research
Industry teaming	Integrate through a prime or platform provider	Reference architectures and integration support	Reference architectures under license

[ENGAGE]

GOVERN AUTONOMY AT RUNTIME.

Burak Oktenli, MBA · Principal Investigator

Independent Researcher, AUTHREX Research Program, Washington, DC

authrex.systems · burakoktenli.com · info@burakoktenli.com · ORCID 0009-0001-8573-1667

TECHNICAL EXCHANGE → SCENARIO MAPPING → SOURCE ACCESS UNDER INVITED REVIEW → SEEDED CONSOLES →

[DISCLAIMERS AND STATUS LABELS]

This catalog presents public-source, self-authored research artifacts. Every result is self-reported and has not been independently reproduced or technically verified by another party. Evaluation is simulation on synthetic data unless a physical test is specifically identified; none is. Hardware references are design specifications; no platform has been built. Formal results are bounded and model-relative; model checking of a configured finite model is not a proof of any implementation or of the physical system. Provisional patent applications are unexamined filings that confer no enforceable rights. Standards are mapped by the author or named as design targets; no compliance, certification, accreditation, or authorization to operate is claimed or implied.

The catalog does not represent official government endorsement, sponsorship, evaluation (beyond the AFRL memorandum described inside), or affiliation. It is not a government publication and carries no DoD distribution statement. Department of Defense is the statutory name used here. Distribution: public. License: CC BY 4.0.

Best read alongside the per-artifact deposits, which contain the full bills of materials, specifications, simulation data, TLA+ modules, and source code, and alongside the correction log, which is the record of what this edition changed and what remains pending.

[STANDING OFFER]

A 30 to 60 minute technical exchange; a written mapping of a government scenario to the authority lifecycle; source access under invited review; government-defined scenarios on the seeded consoles. No mechanism is stated to be available or guaranteed.