



[TECHNICAL PROJECT CATALOG · VERSION 3.2 · JULY 28, 2026]

TECHNICAL PROJECT CATALOG & CAPABILITY OVERVIEW

Authority lifecycle governance infrastructure for safe autonomous systems. Seven governance architectures, twelve hardware platforms, twenty-two reproducible simulations, and a public research corpus, documented at their actual standing for technical review.

7

GOVERNANCE
ARCHITECTURES

12

HARDWARE PLATFORMS

22

REPRODUCIBLE
SIMULATIONS

1

PEER-REVIEWED
ARTICLE

7

PROVISIONAL
PATENTS

9

APPLICATION
DOMAINS

INDEPENDENT RESEARCH PORTFOLIO. AUTHORED BY BURAK OKTENLI UNDER THE AUTHORITY & ARCHITECTURE IMPRINT. NOT A GOVERNMENT PUBLICATION; NOT AFFILIATED WITH, ENDORSED BY, OR PRODUCED FOR ANY GOVERNMENT AGENCY.

DOCUMENT CONTROL AND DISTRIBUTION

Document Title	Technical Project Catalog and Capability Overview
Program	AUTHREX (Authority Lifecycle Governance) research program
Author	Burak Oktenli, Independent Researcher, AI Governance and Safety-Critical Autonomous Systems
Imprint	Authority & Architecture (imprint of record)
Version / Date	Version 3.2, July 28, 2026
Revision note	Revisions 3.1 and 3.2 record the first peer-reviewed journal article (Springer Nature, 2026) and the AFRL BAA evaluation record, add the three governor consoles and the AUTHREX-TEAM and AUTHREX-REDLINE prototypes, and update counts; see the July 28, 2026 addendum.
Identifiers	ORCID 0009-0001-8573-1667 · burakoktenli.com · authrex.systems
Distribution	Public. Open distribution under Creative Commons Attribution 4.0 (CC BY 4.0).
Purpose	To present the author's systems, hardware platforms, simulations, and publications in a structured technical format suitable for review by defense, aerospace, infrastructure, and academic readers.

Notice. This is an independent research portfolio. It is not a government document and is not affiliated with, endorsed by, or produced for the U.S. Department of Defense or any other government agency. All maturity levels, validation states, and limitations are stated plainly so that a technical reviewer can assess the work on its actual standing. Capabilities described at low technology readiness levels are research artifacts, not fielded systems.

Reading note. Technology readiness level (TRL) follows the standard nine-point scale. Open-access deposits (Zenodo) and working papers (SSRN) are public disclosures with registered identifiers; they are not described as peer-reviewed. Provisional patent applications are unexamined filings, not granted patents.

[TABLE OF CONTENTS]

1	Executive Overview	3
2	System Boundary & Readiness Statement	4
3	Capability Areas & Application Domains	5
4	AUTHREX Integrated Authority Lifecycle	6
5	Governance Architecture Register	7
6	Framework Profiles	8
7	BLADE Hardware Governance Platforms	10
8	Flagship Platforms: EDGE, AV, SPACE	11
9	Platform Profiles & Register	14
10	Software Reference Architectures	17
11	Simulation & Modeling Systems	19
12	Publications, Patents & Reference Works	21
13	Technical Readiness Summary	22
14	Evidence & Documentation Matrix	23
15	Mission Alignment Matrix	24
16	Risk & Limitation Assessment	25
17	Roadmap & Priority Improvement Plan	26
18	Summary, Engagement & Glossary	27

EXECUTIVE OVERVIEW

This catalog documents the technical work of Burak Oktenli, an independent researcher in AI governance and safety-critical autonomous systems. The work is conducted under the AUTHREX research program and published through the Authority & Architecture imprint.

[THE PROBLEM]

Autonomous systems are increasingly asked to act faster than a human can supervise, on sensor data that may be degraded or deliberately corrupted, and no standardized architecture governs how authority is granted, constrained, and recovered as conditions change. Policy frameworks such as the NIST AI Risk Management Framework establish principles but do not enforce them in hardware. DoD Directive 3000.09 requires appropriate human judgment over the use of force but does not specify a technical enforcement mechanism. Planning-layer safety systems constrain behavior but do not govern whether a command reaches an actuator. This work proposes an architecture for that missing enforcement layer.

[THE CONTRIBUTION]

Authority itself is treated as an engineered lifecycle. Sensor trust is computed continuously, converted into a graded authority level, screened for adversarial deception, reconciled across multiple agents, gated by a deliberation window before irreversible action, and recovered deterministically to a safe state when trust collapses. The same pipeline applies across nine application domains without retraining the underlying autonomy.

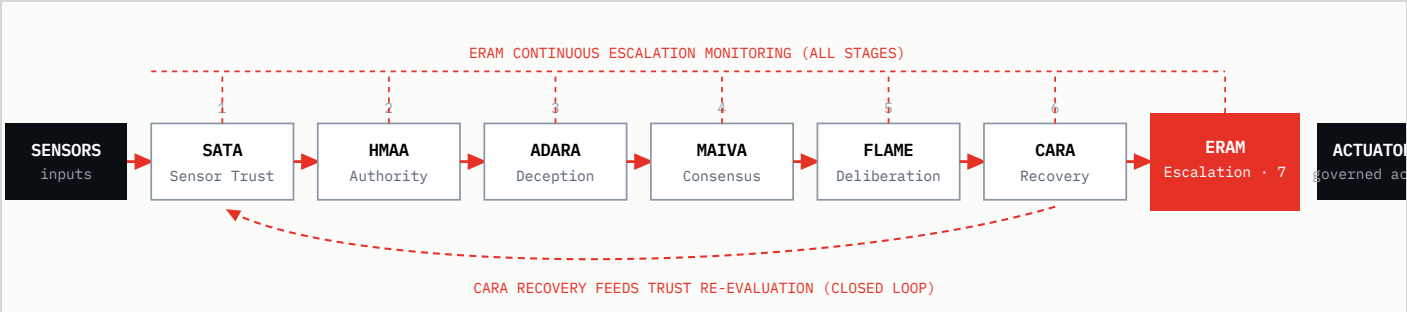


Figure 1. The AUTHREX seven-stage authority lifecycle. Sensor inputs enter at left; governed action exits at right. ERAM monitors all stages continuously, and CARA recovery feeds back into trust re-evaluation, forming a closed loop rather than a one-shot pipeline.

[WHAT THIS CATALOG CONTAINS]

Seven governance architectures, twelve hardware research platforms, six software reference architectures, twenty-two browser-based simulations, the public research corpus (registry pending), seven provisional patent applications, one peer-reviewed journal article, and a ten-volume technical reference series. These artifacts share one architectural foundation and are documented individually with maturity, evidence, and domain relevance consolidated in the matrices of Sections 13 through 15.

SYSTEM BOUNDARY AND READINESS STATEMENT

The scope of the work is specific. AUTHREX governs the path between the autonomy software and the actuator; it does not replace the autonomy and does not make mission or tasking decisions. Figure 2 shows that boundary, including the hardware interlock and the operator, so a reviewer can see exactly what the architecture controls and what remains outside it.

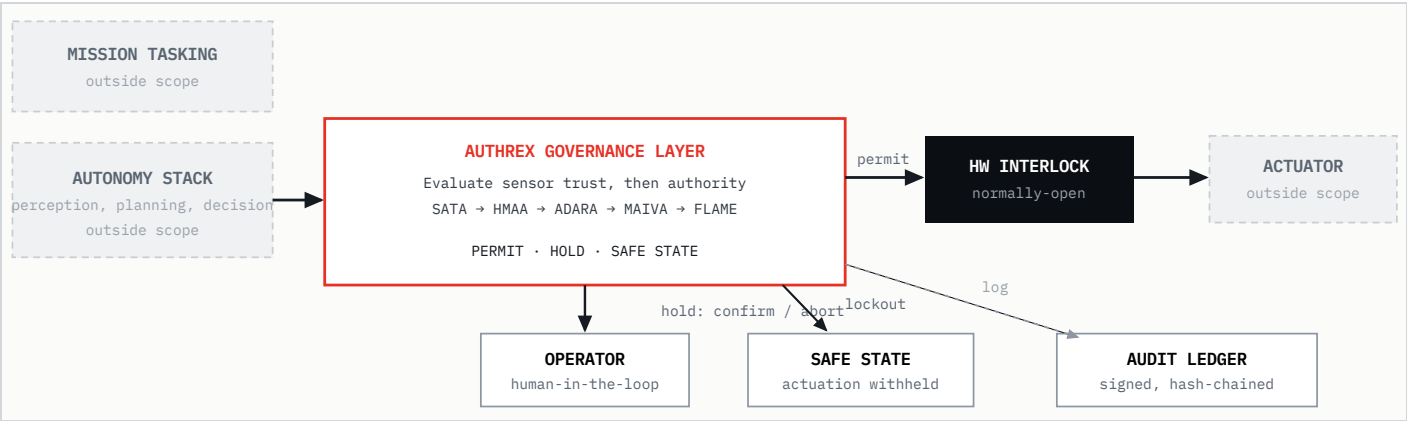


Figure 2. AUTHREX system boundary. The autonomy stack proposes an action; the governance layer evaluates trust and authority and either permits, holds for operator confirmation, or commands a safe state; a hardware interlock sits between the decision and the actuator; every decision is written to an audit ledger. Autonomy logic, mission tasking, and the physical actuator remain outside the governance scope.

[WHY THE WORK MAY BE RELEVANT]

The architecture maps onto stated public priorities: autonomous-systems human-control requirements (DoD Directive 3000.09); trustworthy AI in critical infrastructure (NIST AI RMF); automated-vehicle safety and false vehicle-control commands (NHTSA framework and the SELF DRIVE Act of 2026, H.R. 7390); orbital autonomy under signal delay (NASA SBIR EXPAND.3.S26B and Space Policy Directive 5); safe adoption of agentic AI (CISA, NSA, and Five Eyes guidance, May 2026); counter-unmanned-aircraft authority (Executive Order 14305 and the FY26 NDAA); and operational-technology security (NIST SP 800-82, ISA/IEC 62443, NERC CIP). The work is structured so that future evaluation can be mapped against DO-178C and DO-333 for airborne software, MIL-STD-882E for system safety, and ISO 26262 for automotive.

[HONEST READINESS STATEMENT]

The program is at an early technology readiness level, generally self-assessed TRL 2 to 4. Simulations demonstrate the governance logic with seeded, reproducible runs, and the HMAA authority state machine is formally verified in TLA+ (23,748 distinct reachable states at depth 9; of 8 stated properties, 6 hold and 2 are vacuous at this bound). The hardware platforms are specified to the bill-of-materials level but are not yet built. Independent peer review and red-team evaluation are planned, not complete. The concept addresses a defined governance gap and the architecture is internally consistent; it is not yet proven, deployment-ready infrastructure, and this catalog does not present it as such.

CAPABILITY AREAS AND APPLICATION DOMAINS

The portfolio organizes into the capability areas below. Only areas genuinely supported by the work are listed; each is cross-referenced in the readiness, evidence, and alignment matrices.

CAPABILITY AREA	SYSTEMS IN THIS CATALOG	CURRENT MATURITY
AI and Autonomy Governance	AUTHREX pipeline; SATA, HMAA, ADARA, MAIVA, FLAME, CARA, ERAM	Research / simulation-demonstrated
Human-Machine Systems	HMAA; FLAME deliberation gating	Research / simulation
Defense Technology	BLADE-EDGE, BLADE-CUAS, BLADE-SWARM, BLADE-AGENT-HSM	Design complete / BOM-specified
Cyber and Security	AUTHREX-AGENT, AUTHREX-AGENT-CYBER, AUTHREX-ASSURE	Reference architecture
Resilient Infrastructure	BLADE-INFRA, BLADE-INFRA-OT, AUTHREX-ICS-GATE	Design / reference
Space and Aerospace	BLADE-SPACE, AUTHREX-SPACECYBER	Preliminary design (self-assessed TRL 2-3)
Economic and Supply-Chain Security	BLADE-FINANCE	Simulation (synthetic data only)
Simulation Systems	ERAM Strategic Command, JADC2 Theater, MDO Testbed, and others total) (22	Operational research tools
Technical Research and Publications	Public works, provisional filings, 10-volume reference series	Published / disclosed

[NINE APPLICATION DOMAINS · ONE GOVERNANCE LOGIC]

The same seven-stage governance logic is instantiated against nine distinct application domains with domain-specific sensors, standards, and hardware. This breadth is the basis for describing AUTHREX as a cross-domain governance architecture rather than a single-domain product.

AIRCRAFT Flight control integrity	GROUND VEHICLE Drive-by-wire safety	MARITIME USV/UUV Surface and subsurface
DEFENSE UAS Authority handoff	POWER GRID ICS / SCADA control	SPACE / ORBITAL Onboard autonomy
AGENTIC AI Tool-use governance	CYBER-DEFENSE Autonomous patching	CRITICAL OT IT / OT boundary

AI AND AUTONOMY: AUTHORITY GOVERNANCE

SYSTEM · UNIFIED GOVERNANCE PIPELINE

AUTHREX Integrated Authority Lifecycle

RESEARCH ARCHITECTURE · self-assessed TRL 3-4



[PROJECT STATUS]

Research architecture integrating seven published frameworks. Demonstrated in seeded, reproducible simulation runs; formally verified in part. Not deployed.

[SUMMARY]

AUTHREX is a governance layer that sits between an autonomous system’s decision logic and its actuators. It does not make the system more intelligent. It determines, in real time, whether the system is permitted to act, under what authority, and what happens when trust degrades. The intelligence stays in the autonomy; the authority is governed.

[TECHNICAL DESCRIPTION]

SATA computes a continuous sensor-trust scalar using weighted Dempster-Shafer fusion. HMAA converts trust and context into a graded authority level via a finite-state machine. ADARA reduces authority when deception is likely. MAIVA reconciles authority across agents using Byzantine fault-tolerant voting (3f+1). FLAME imposes a deliberation window scaled to consequence before irreversible actions. CARA executes deterministic recovery to safe states on lockout. ERAM monitors all stages and assesses escalation risk. Enforcement is designed to terminate at a hardware boundary, a normally-open relay between decision logic and actuator, so that compromise of the software stack does not by itself permit actuation.

[KEY CAPABILITIES]

- Graded, non-binary authority proportional to real-time trust
- Continuous per-source sensor-trust assessment (Dempster-Shafer)
- Adversarial deception screening before authority is committed
- Byzantine-resilient multi-agent consensus (f faults in 3f+1)
- Consequence-scaled deliberation gating; deterministic recovery
- Cryptographically signed decision audit chain (ECDSA P-256)

[INTENDED USE CASES]

- Defense autonomy authority gating (mass autonomy, CCA)
- Automotive driver-assist handoff on an ASIL-D pathway
- ICS / SCADA and IT/OT boundary command adjudication
- Orbital autonomy under ground-loop signal delay
- Agentic-AI tool-use containment
- Counter-UAS detection-to-authority handoff

VALIDATION: 2,800+ SIM RUNS ON REPRESENTATIVE PLATFORMS, NO UNSAFE ACTIONS IN SIMULATION NO HIL · NO INDEPENDENT RED-TEAM YET
· HMAA TLA+ VERIFIED

SEVEN FRAMEWORKS, ONE LIFECYCLE

CODE	NAME	FUNCTION	IP / DISCLOSURE	VALIDATION TO DATE
SATA	Sensor & Actor Trust Assessment	Continuous sensor-trust scalar via weighted Dempster-Shafer fusion	U.S. Provisional 64/002,453	Sim + TPM-anchor design
HMAA	Human-Machine Authority Allocation	Graded four-level authority (A3-A0) with asymmetric hysteresis	U.S. Provisional 63/999,105	TLA+ verified (23,748 states)
ADARA	Adversarial Deception-Aware Reasoning	Authority adjustment by deception prior P(adversarial)	Zenodo deposit	Simulation
MAIVA	Multi-Agent Integrity Voting	Byzantine-resilient consensus (3f+1), CUSUM detection	Zenodo deposit (MIT)	TLA+ spec, 37 self-tests
FLAME	Forced Latency for Authority-Managed Escalation	Consequence-scaled deliberation window; 5-state circuit breaker	U.S. Provisional 64/005,607	Simulation
CARA	Controlled Authority Recovery Architecture	Deterministic GREP-phase recovery; terminal policy gate	U.S. Provisional 64/000,170	10M-iteration verification
ERAM	Escalation Risk Assessment Model	Escalation-risk quantification for AI-enabled command and control	SSRN working paper	Sim (600 Monte Carlo runs)

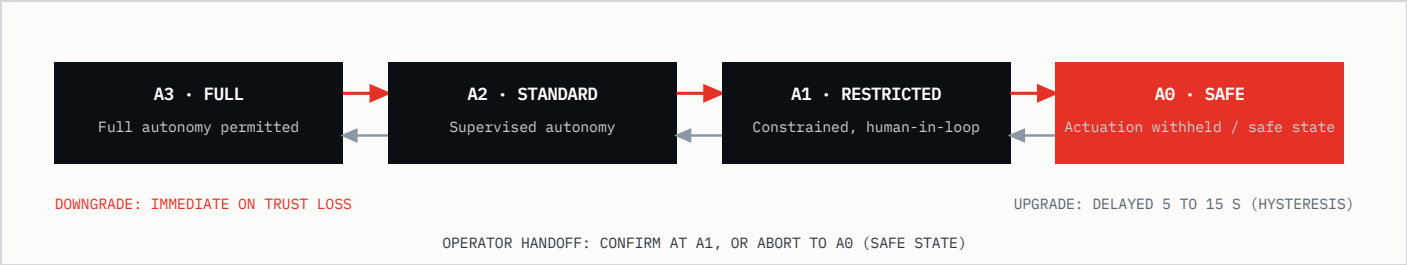


Figure 3. HMAA graded authority levels with asymmetric hysteresis and the operator handoff path. Downgrade is immediate on loss of trust; upgrade is delayed to prevent oscillation. At the restricted level the operator confirms or aborts; the lowest level withholds actuation and commands a safe state.

Reviewer note: the novel element across the register is the treatment of authority as an engineered, trust-proportional lifecycle with formal and hardware anchoring, rather than a binary safety override. Current standing is pre-hardware and pre-independent-review; assess as research architecture, not fielded capability.

SATA · HMAA · ADARA · MAIVA

SATA · Sensor & Actor Trust AssessmentProv. 64/002,453 (filed 11 Mar 2026) · DOI 10.5281/zenodo.18936251

Foundation trust layer. Produces a continuous trust scalar in the interval zero to one from a hardware-anchored attestation chain, using weighted Dempster-Shafer belief functions across four diagnostics: internal consistency, cross-sensor agreement, temporal stability, and physical plausibility. Every downstream stage consumes this scalar.

HMAA · Human-Machine AuthorityAllocationProv. 63/999,105 (filed 7 Mar 2026) · DOI 10.5281/zenodo.18861653

Converts trust and operational context into a graded, four-level authority state (A3 through A0) with asymmetric hysteresis: authority is reduced immediately on trust loss and restored only after a sustained delay of roughly 5 to 15 seconds. Implemented as a 42-file package with 98 tests and formally verified in TLA+ across 23,748 distinct reachable states at depth 9; of 8 stated properties, 5 invariants and 1 liveness hold and 2 are vacuous at this bound.

ADARA · Adversarial Deception-Aware ReasoningDOI 10.5281/zenodo.19043924

A proactive deception prior that adjusts authority downward in proportion to the probability that inputs have been manipulated. P(adversarial) is computed from input anomalies, temporal correlation, cross-sensor consistency, and Bayesian mission history, with a phantom-fleet module to resist AI-hallucinated or fabricated hostile scenarios.

MAIVA · Multi-Agent Integrity VotingDOI 10.5281/zenodo.19015517 · MIT license

Byzantine-resilient aggregation of authority across multiple agents using a trimmed weighted median resistant to f adversaries in a 3f+1 roster, with CUSUM-augmented anomaly detection and DoDD 3000.09 action-gate classification. Carries a TLA+ specification and 37 self-tests.

FLAME · CARA · ERAM

FLAME · Forced Latency for Authority-Managed Escalation

Prov. 64/005,607 (filed 14 Mar 2026) · DOI 10.5281/zenodo.19015618

Treats strategic latency as an engineered system. Before any irreversible action, FLAME imposes a mandatory deliberation window sized by a dynamic delay function over authority, consequence tier, and domain. A five-state circuit breaker with cryptographically signed transitions enforces the delay. The fail-safe default on timeout is abort, not execute.

CARA · Controlled Authority Recovery Architecture

Prov. 64/000,170 (filed 9 Mar 2026) · DOI 10.5281/zenodo.18917790

Deterministic recovery from authority lockout through GREP phases (Guard, Reduce, Evaluate, Promote) with a terminal non-compensatory policy gate, so that recovery cannot be short-circuited by a single favorable signal. Verified by 10-million-iteration enumeration and a 68-state control-flow analysis.

ERAM · Escalation Risk Assessment Model

SSRN ID 6176802 · working paper

Quantifies decision-time compression and escalation pathways in AI-enabled command and control, modeling how local autonomous actions cascade across domain boundaries. Implemented as the ERAM Strategic Command simulation with 600 Monte Carlo runs and cryptographic provenance.

[Technology Readiness Position]




Figure 4. Technology readiness position. Governance logic is demonstrated in reproducible simulation and, for HMAA, in formal verification (self-assessed TRL 2 to 4). Hardware platforms are specified to the bill-of-materials level but are not yet built. The stated near-term target is TRL 6 through hardware-in-the-loop validation.

BLADE HARDWARE GOVERNANCE PLATFORMS

The BLADE family instantiates the governance pipeline as domain-specific hardware nodes. Each platform is published as a reproducible artifact package with a full bill of materials, interface documentation, and simulation data. All platforms share one pipeline architecture, which is the basis for cross-domain reuse.

HIGH-CONSEQUENCE DEFENSE · REFERENCE HARDWARE

BLADE-EDGE Governance Node

DESIGN COMPLETE · self-assessed TRL 3



[PROJECT STATUS]

Design complete; prototype specification. Bill of materials defined; not built.

[SUMMARY]

A ruggedized edge-computing governance node that hosts the full nine-module governance pipeline for high-consequence defense platforms. The normally-open safety interlock is designed so that a governed actuator cannot act without authorization, even under software compromise.

[TECHNICAL DESCRIPTION]

Dual-redundant NVIDIA Jetson AGX Orin compute paired with dual Xilinx Zynq UltraScale+ FPGAs. Nine-module pipeline (SATA, ADARA, IFF, HMAA, MAIVA, FLAME, CARA, outcome assessment, actuation gate). 72 components, 103 connections, MIL-STD-810G environmental design target, hardwired normally-open safety interlock relay.

[EVIDENCE]

Zenodo deposit DOI 10.5281/zenodo.19177472; full bill of materials, interface control documentation, and nine-module pipeline specification; reference cost approximately 139,000 dollars.

[LIMITATIONS & REVIEWER NOTE]


Highest-cost platform and not built. Actuation-governance claims rest on simulation and design analysis; the MIL-STD-810G rating is a design target, not a test result. The most developed hardware reference design in the portfolio and the basis for cross-domain reuse (roughly 75 percent reuse into BLADE-CUAS). Read as a detailed reference design pending build and test.

BLADE HARDWARE GOVERNANCE PLATFORMS

AUTOMOTIVE DRIVE-BY-WIRE · ISO 26262 ASIL-D PATHWAY

BLADE-AV Governance Node

PUBLISHED · self-assessed TRL 3-4



[PROJECT STATUS]

Published on Zenodo. Design complete and demonstrated in simulation; hardware not built.

[MISSION RELEVANCE]

Addresses a structural gap in driver-assist and automated-driving safety. NHTSA documented 1,429 automated-vehicle incidents between 2021 and 2025, and the SELF DRIVE Act of 2026 (H.R. 7390) requires cybersecurity plans for false vehicle-control commands. BLADE-AV is a reference architecture for hardware-enforced command governance.

[TECHNICAL DESCRIPTION]

Nine-module pipeline on Jetson AGX Orin plus Zynq UltraScale+. Three-leg redundant KILOVAC LEV200 fail-safe relay. 62 components. Trust-driven forced driver handoff and brake-to-stop-in-lane degradation. ISO 26262 ASIL-D and SAE J3016 Level 4 design targets.

[EVIDENCE]

Zenodo deposit DOI 10.5281/zenodo.19232130; 1,200 simulation runs with no observed unsafe authority transitions recorded in simulation; reference cost approximately 16,300 dollars. A co-authored journal manuscript is in preparation with Dr. Victor Lian (Lynn University).

[VALIDATION & LIMITATIONS]

1,200 reproducible simulation runs; cross-domain portability demonstrated in simulation against BLADE-EDGE. No vehicle integration; ASIL-D is a target pathway, not a certification; the journal manuscript is in preparation and is not a published, peer-reviewed paper.

[REVIEWER NOTES]


Strong simulation evidence and a clear qualification pathway. Published as an artifact but remains pre-hardware. Development path: bench integration of the fail-safe relay, hardware-in-the-loop with a drive-by-wire rig, completion of the co-authored manuscript.

BLADE HARDWARE GOVERNANCE PLATFORMS

ORBITAL LEO AUTONOMY · NASA SBIR TOPIC MAPPING

BLADE-SPACE Governance Node

PRELIMINARY DESIGN · self-assessed TRL 2-3



[PROJECT STATUS]

Preliminary design phase. Engineering package complete; verification and validation specified but not executed.

[MISSION RELEVANCE]

Potentially relevant to NASA SBIR subtopic EXPAND.3.S26B, where onboard authority governance supports autonomous onboard health management for small spacecraft and distributed systems, and to Space Policy Directive 5. Addresses autonomous conjunction-avoidance and payload-firing decisions made on stale or spoofed orbital data.

[TECHNICAL DESCRIPTION]

Microchip RTG4 radiation-tolerant FPGA plus Aitech S-A1760 Venus single-board computer in hot redundancy with sub-200-millisecond failover. Nine-stage pipeline. ECDSA P-256 audit chain anchored in a radiation-tolerant TPM. Three-fault-tolerant payload and thruster firing interlock. ADARA multi-constellation GNSS spoofing detection. 91 components, 6U-plus SmallSat payload module, 30 krad total ionizing dose, five-year design life.

[EVIDENCE]

Zenodo deposit DOI 10.5281/zenodo.20183269; 15-document engineering package; reference cost approximately 505,000 dollars.

[VALIDATION & LIMITATIONS]

Preliminary design with a 15-document package. The verification and validation campaign is specified but not executed; no platform-level simulation campaign completed. The earliest-maturity hardware platform.

[REVIEWER NOTES]

A thorough preliminary design and the least mature platform. Treat as a design study with a defined but unexecuted validation plan. Development path: execute the specified V&V, build the orbital simulation campaign, pursue NASA SBIR alignment.

ROVER · UAV · MARITIME · INFRA

Rover Testbed

Design complete, build in progress · DOI 10.5281/zenodo.19143190 · ~\$484

A low-cost reference testbed for authority-governed autonomy. Dual compute (Raspberry Pi 5 and ESP32) running the SATA-HMAA-CARA pipeline; 37 verified components, 76 electrical connections, 7 fault scenarios. 350 simulation runs recorded no observed unsafe authority transitions in simulation; the HMAA state machine is TLA+ verified.

Robotics testbed · ROS 2 · self-assessed TRL 3-4

HMAA-UAV Platform

Design complete, build in progress · DOI 10.5281/zenodo.19128769 · ~\$4,200

Trust-governed flight autonomy for contested environments. Cube Orange+ flight controller with an NVIDIA Jetson Orin NX companion and a MAVLink hardware-in-the-loop bridge; 52 components. 250 simulation runs across 5 adversarial scenarios, with a Monte Carlo campaign.

UAS testbed · ArduPilot · self-assessed TRL 3-4

BLADE-MARITIME

Published · DOI 10.5281/zenodo.19246785 · ~\$43K

Maritime surveillance governance with hydroacoustic sonar, magnetic anomaly detection, and AIS spoofing detection. Four maritime mathematical extensions include sea-state authority damping and acoustic-delay-aware Byzantine consensus. 84 components, IP68 enclosure, MIL-STD-810G and MIL-STD-461G CE102 design targets.

Maritime (surface / subsurface) · self-assessed TRL 3 (design complete)

BLADE-INFRA

Published · DOI 10.5281/zenodo.19277887 · ~\$11.6K

Critical-infrastructure protection node for ICS and SCADA, power-grid monitoring, water treatment, and pipeline operations. IEC 61850 GOOSE, Modbus TCP/RTU, and PROFINET IO; Pilz PNOZ S7.1 SIL-3 safety relay; 92 components; SIL 3, NERC CIP, and FIPS 140-2 Level 3 design targets.

Critical infrastructure · self-assessed TRL 3 (design complete)

CUAS · AGENT-HSM · SWARM · INFRA-OT · FINANCE

BLADE - CUAS

DOI 10.5281/zenodo.20299604 · ~\$43.5K · self-assessed TRL 2-3 hw / 3-4 sim

Passive counter-UAS governance placed between commercial detection sensors and authorized operators. Four-level HMAA federal-to-SLTT authority handoff; Dempster-Shafer consensus across radar, RF, EO/IR, Remote ID, and LIDAR; ECDSA P-256 tamper-evident evidence record. Aligned with Executive Order 14305 and the FY26 NDAA Safer Skies Act; approximately 75 percent design reuse from BLADE-EDGE.

Counter-UAS · sixth domain instantiation

BLADE - AGENT - HSM

DOI 10.5281/zenodo.20299821 · ~\$199 · self-assessed TRL 2-3 silicon / 3-4 emulator

Tamper-evident hardware root of trust for autonomous agents and the hardware companion to the AUTHREX-AGENT shim. Non-exportable ECDSA P-256/P-384 keys in an NXP EdgeLock SE051 (CC EAL6+); four-level authority state in an Infineon SLB 9670 TPM 2.0; HKDF per-tool tokens; multi-modal tamper cascade that zeroizes keys. Five-opcode 64-byte ABI in USB-A and M.2 Key-E form factors; verified by an adversarial emulator (275 of 275 deterministic checks).

Agentic-AI root of trust · first HSM in the family

BLADE - SWARM

DOI 10.5281/zenodo.20351198 · ~\$1,333/node · self-assessed TRL 3-4 sim / 2 testbed

Authority governance for attritable multi-agent swarms. Byzantine-fault-tolerant two-phase consensus gated by SATA, HMAA, and MAIVA across N = 10, 50, and 500 agents, tolerating $f = (N-1)/3$ compromised agents per quorum with a quorum-intersection bound and safe-halt-by-default under denied RF. Per-node ECDSA P-256 root of trust and a hash-chained audit ledger; TLA+ verified (5 safety, 3 liveness). Governs decision authority and audit only.

Swarm autonomy · multi-agent extension of HMAA-UAV

BLADE - INFRA - OT

DOI 10.5281/zenodo.20342067 · ~1U appliance · self-assessed TRL 2-3 hw / 3-4 sim

A fail-closed, bump-in-the-wire governance appliance at the IT/OT segmentation boundary. Each cross-boundary command is adjudicated to propagate, hold for deliberation, or isolate across four OT regimes (nominal, elevated, lockdown, safe-halt); malformed input fails closed. Xilinx Kria K26 governance plane and x86 network plane; 48 BOM line items; SHA-256 tamper-evident audit ledger.

Operational technology · OT companion to BLADE-INFRA

BLADE - FINANCE

DOI 10.5281/zenodo.20374692 · ~\$9,228 · self-assessed TRL 3-4 sim / 2 hw

Authority arbitration for financial-sector AI decision systems under the U.S. Treasury Financial Services AI Risk Management Framework. An eight-stage pipeline routes each transaction to autonomous clearance, supervised review, elevated confirmation, or manual hold; retrospective swarm-review recovers coordinated low-and-slow activity; SHA-256 evidence chain; YubiHSM 2 in a FIPS 140-2 Level 3 enclosure; 36 components. Synthetic data only; not deployed in any institution.

Economic security · first platform in the domain

TWELVE HARDWARE RESEARCH PLATFORMS

All standards listed are design targets and research mappings, not accreditations or audit findings. Component counts marked with a hyphen are not stated as a single figure in the source artifact. TRL pairs denote hardware / simulation maturity.

PLATFORM	DOMAIN	COMP.	REF. BOM	TRL	STANDARDS ALIGNMENT (DESIGN TARGETS)	ZENODO DOI
Rover Testbed	Robotics testbed	37	~\$484	3-4	ROS 2; TLA+ verified	10.5281/zenodo.19143190
UAV (HMAA-UAV)	UAS testbed	52	~\$4,200	3-4	MAVLink/HIL; ArduPilot	10.5281/zenodo.19128769
BLADE-EDGE	High-consequence defense	72	~\$139K	3 (design)	MIL-STD-810G	10.5281/zenodo.19177472
BLADE-AV	Automotive drive-by-wire	62	~\$16.3K	3-4	ISO 26262 ASIL-D; SAE J3016 L4	10.5281/zenodo.19232130
BLADE-MARITIME	Maritime surveillance	84	~\$43K	3 (design)	IP68; MIL-STD-810G / 461G	10.5281/zenodo.19246785
BLADE-INFRA	Critical infrastructure (ICS/SCADA)	92	~\$11.6K	3 (design)	SIL 3; NERC CIP; FIPS 140-2	10.5281/zenodo.19277887
BLADE-SPACE	Orbital LEO autonomy	91	~\$505K	2-3 (prelim)	NASA EXPAND.3.S26B; SPD-5	10.5281/zenodo.20183269
BLADE-CUAS	Counter-UAS	-	~\$43.5K	2-3 / 3-4	EO 14305; FY26 NDAA Safer Skies	10.5281/zenodo.20299604
BLADE-AGENT-HSM	Agentic-AI root of trust	-	~\$199	2-3 / 3-4	CC EAL6+; TPM 2.0; Five Eyes guidance	10.5281/zenodo.20299821
BLADE-SWARM	Attributable swarm autonomy	-	~\$1,333/node	3-4 / 2	DoDD 3000.09; NIST AI RMF	10.5281/zenodo.20351198
BLADE-INFRA-OT	IT/OT boundary governance	48	~1U unit	2-3 / 3-4	NIST SP 800-82; IEC 62443; NERC CIP	10.5281/zenodo.20342067
BLADE-FINANCE	Financial-sector AI	36	~\$9,228	3-4 / 2	Treasury FS AI RMF; EO 14179	10.5281/zenodo.20374692

[SIMULATED CAMPAIGN RESULTS]

Across representative platforms, more than 2,800 simulation runs recorded no observed unsafe authority transitions in simulation (Rover 350, UAV 250, BLADE-AV 1,200, plus per-framework campaigns). HMAA authority computation completes in under 12 ms (simulated); MAIVA tolerates $f < n/3$ Byzantine nodes. Simulated performance characteristics are research artifacts, not fielded-system measurements.


SIX SOFTWARE REFERENCE ARCHITECTURES

Six software reference architectures carry the same authority pipeline into settings that do not require a dedicated hardware node. All are single-author research and are not deployed in any production environment.

AGENTIC-AI AUTHORITY SHIM · SOFTWARE-ONLY

AUTHREX-AGENT

REFERENCE ARCHITECTURE · self-assessed TRL 3-4



[PROJECT STATUS]

Reference architecture with a working simulator. Single-author research; not deployed.

[MISSION RELEVANCE]

Aligned with CISA, NSA, and Five Eyes guidance Careful Adoption of Agentic AI Services (1 May 2026) and FY26 NDAA Sections 1513 and 6601. Tool misuse and credential exposure are named risk categories for agentic AI.

[TECHNICAL DESCRIPTION]

Four-level (T3 to T0) HMAA authority gating; per-tool HKDF authorization tokens; sub-agent spawn quorum; SATA input-trust scoring that collapses on credential patterns or prompt-injection signatures; FLAME bounded deliberation with a fail-safe abort on timeout; hash-chained audit ledger. Hardware-anchored by the BLADE-AGENT-HSM root of trust when present.

[EVIDENCE]

A 15-section technical specification with a standards-alignment matrix, an SDK integration guide, and three reference use cases; verified against the AUTHREX-AGENT simulator; documented hardware companion (DOI 10.5281/zenodo.20299821).

[LIMITATIONS & REVIEWER NOTE]

Software-only containment can be bypassed without the hardware root of trust; no production deployment; no independent evaluation. The most directly implementable near-term software application of the pipeline and the closest fit to current federal agentic-AI guidance.

ASSURE · ICS-GATE · AGENT-CYBER · SPACECYBER · SANDBOX

AUTHREX-ASSURE

Pre-deployment assurance gate · reference architecture (self-assessed TRL 3-4)

An assurance gate a system must clear before it is permitted to act, addressing the validation gap named in the 2026 National Cybersecurity Strategy. Aligned with NDAA Section 1533. Not deployed.

AUTHREX-ICS-GATE

OT authority governance at the IT/OT boundary · reference architecture

Gates AI actions across the IT/OT boundary for critical infrastructure; it authorizes deterministic safety logic and does not make the safety decision itself. Companion to BLADE-INFRA-OT. Aligned with CISA and NSA AI-in-OT principles, NIST SP 800-82, ISA/IEC 62443, and NERC CIP. Not deployed.

AUTHREX-AGENT-CYBER

Autonomous cyber-defense authority · governance only, no offensive function

Constrains what a defensive cyber-reasoning agent may do and when, for example whether it may patch a live OT controller. It treats the agent as a black box and performs no vulnerability discovery itself. Aligned with Five Eyes agentic-AI guidance, DARPA AIxCC, and NDAA Section 1513. Not deployed.

AUTHREX-SPACECYBER

Onboard authority for orbital autonomy · reference architecture

Gates autonomous spacecraft actions under intermittent ground contact. Companion to BLADE-SPACE. Aligned with NASA SBIR EXPAND.3.S26B and Space Policy Directive 5. Not deployed.

AUTHREX-SANDBOX

Test-and-evaluation authority governance · reference architecture

Governs what an AI under evaluation is permitted to do inside a test environment. Aligned with NDAA Sections 1533 and 1534. Not deployed.

SYSTEM	FUNCTION	STANDARDS ALIGNMENT	MATURITY (SELF-ASSESSED TRL)	STATUS
AUTHREX-AGENT	Authority-lifecycle shim for agentic AI	CISA/NSA/Five Eyes; NDAA 1513, 6601	3-4	Spec + simulator; not deployed
AUTHREX-ASSURE	Pre-deployment assurance gate	2026 Nat'l Cyber Strategy; NDAA 1533	3-4	Reference; not deployed
AUTHREX-ICS-GATE	OT authority at the IT/OT boundary	CISA/NSA AI-in-OT; NIST 800-82; IEC 62443; NERC CIP	3-4	Reference; not deployed
AUTHREX-AGENT-CYBER	Autonomous cyber-defense authority (governance only)	Five Eyes; DARPA AIxCC; NDAA 1513	3-4	Reference; not deployed
AUTHREX-SPACECYBER	Onboard authority for orbital autonomy	NASA EXPAND.3.S26B; SPD-5	3-4	Reference; not deployed
AUTHREX-SANDBOX	Test-and-evaluation authority governance	NDAA 1533, 1534	3-4	Reference; not deployed

REPRODUCIBLE RESEARCH TOOLS

The simulations run entirely client-side with a seeded pseudo-random number generator for bit-exact reproducibility, which lets any reviewer independently verify a result. They are the primary current evidence for the governance logic.



Figure 5. Simulation validation workflow. A fixed seed and a scenario definition drive each run; the governed pipeline executes; an unsafe-action check is applied; the run is written to a hash-chained audit log and a result hash is recorded.

ERAM Strategic Command

OPERATIONAL RESEARCH TOOL

Operational browser-based simulation. The underlying ERAM framework is a working paper on SSRN (ID 6176802).

Six scenarios, 600 Monte Carlo runs, formal property verification, and Merkle provenance chains. Runs fully client-side with a seeded PRNG. Integrates ERAM with SATA, FLAME, and CARA to model how authority integrity degrades and how local autonomous actions cascade across domain boundaries in joint all-domain command and control.

Reproducible simulation outputs; model-based with documented parameter assumptions; not validated against operational data. The framework is a working paper, not a peer-reviewed publication.

A reproducible analytical tool, useful as a demonstration and study environment rather than an operational decision aid. Development path: peer-reviewed submission, sensitivity analysis, integration with the JADC2 theater simulation.

SEVEN FLAGSHIP STANDALONE SIMULATIONS

The portfolio comprises 22 simulations in total: these seven, the three governor consoles and the AUTHREX-TEAM and AUTHREX-REDLINE prototypes recorded in the July 2026 addendum, plus per-framework subsystem demonstrations and per-platform validation runs. All are launchable as interactive consoles on [authrex.systems](#).

SIMULATION	FRAMEWORKS EXERCISED	METHOD	SIZE
ERAM Strategic Command	ERAM, SATA, FLAME, CARA	6 scenarios, 600 Monte Carlo, formal properties, Merkle	86 KB
APEX JADC2 Theater	All seven	Full pipeline, PBFT, Merkle audit trail	28 KB
MDO Digital Twin	HMAA, CARA, MAIVA	Air / maritime / ground domains, Byzantine fault isolation	21 KB
Tactical COP (blue-on-blue)	HMAA, CARA, SATA, ADARA	Blue-on-blue interdiction via common operating picture	21 KB
AUTHREX OS	All seven	Distributed Web Worker nodes, COP, analytics, Merkle	-
Distributed Kernel	MAIVA, CARA	Actor-model mesh, PBFT consensus, emergent interlock	-
Tactical Core	HMAA, CARA	Zero-dependency kinematic engine, safe-abort handling	-

[WHAT THE SIMULATIONS DEMONSTRATE, AND WHAT THEY DO NOT]

The simulations demonstrate that the governance logic behaves as specified under seeded scenario conditions: authority degrades with trust, deliberation windows hold irreversible actions, Byzantine minorities are outvoted, and recovery follows the defined staircase. They do not prove field performance, timing on target hardware, or resistance to a live adversary; those claims await hardware-in-the-loop testing and independent red-team evaluation.

PUBLIC DISCLOSURE CORPUS

[OPEN-ACCESS CORPUS]

The public corpus comprises open-access Zenodo deposits with registered DOIs and SSRN working papers; a record-level registry is pending, so no aggregate total is stated. One peer-reviewed journal article is published: Journal of Hardware and Systems Security (Springer Nature) 10, 17 (2026), DOI 10.1007/s41635-026-00190-4. Deposits carry MIT and CC BY 4.0 licenses; claims are verifiable via USPTO, Zenodo, SSRN, and ORCID.

[PROVISIONAL PATENT APPLICATIONS]

ARCHITECTURE	APPLICATION NO.	FILED	RECEIPT	DEPOSIT	STATUS
HMAA	63/999,105	7 Mar 2026	74759595	zenodo.18861653	Provisional; pending USPTO action
CARA	64/000,170	9 Mar 2026	74767602	zenodo.18917790	Provisional; pending USPTO action
SATA	64/002,453	11 Mar 2026	74808459	zenodo.18936251	Provisional; pending USPTO action
FLAME	64/005,607	14 Mar 2026	74858888	zenodo.19015618	Provisional; pending USPTO action
ADARA	64/110,218	13 Jul 2026	pending	zenodo.19043924	Provisional; receipt pending
MAIVA	64/110,221	13 Jul 2026	pending	zenodo.19015517	Provisional; receipt pending
ERAM	64/110,225	13 Jul 2026	pending	ssrn.6176802	Provisional; receipt pending

Provisional applications are unexamined filings; the twelve-month window to file non-provisional (utility) applications runs to approximately March 2027 for the March 2026 filings and July 2027 for the July 2026 filings.

[TECHNICAL REFERENCE SERIES]

A ten-volume technical reference series is published under the Authority & Architecture imprint across three sub-series: The Authority Equation (3 volumes), The Authority Discipline (4 volumes), and Autonomous Authority (3 volumes). ISBNs are registered through Bowker, and The Authority Equation Volume I carries Library of Congress Control Number 2026912260. The series is distributed in print and electronic formats.

[EXTERNAL RECOGNITION]

NHTSA cited the author by name in Federal Register 91 FR 30789, footnote 10. Commentary and analysis have appeared in Modern War Institute at West Point, RUSI, The Defense Post, RealClearDefense, The Space Review, and Military AI, with editorial selection by UPI's Korea Regional Review and the Parliamentary Assembly of the Mediterranean's AI observatory. The full record is maintained at authrex.systems/news.html.

TECHNICAL READINESS SUMMARY

Ratings are stated conservatively. Technical maturity uses the standard TRL scale. Validation describes the strongest evidence that currently exists, and review readiness reflects how prepared an item is for external technical review in its present state.

ITEM	STATUS	MATURITY (SELF- ASSESSED TRL)	DOCS	VALIDATION (STRONGEST CURRENT EVIDENCE)	REVIEW READINESS
AUTHREX integrated pipeline	Research architecture	3-4	High	Simulation; HMAA formally verified	Moderate
SATA, HMAA, CARA, FLAME	Provisional applications filed	3-4	High	Simulation; HMAA TLA+ verified	Moderate
MAIVA, ADARA, ERAM	Open-access / SSRN	3-4	High	Simulation	Moderate
BLADE-EDGE	Design complete	3	High	Simulation + design analysis	Moderate
BLADE-AV	Published (Zenodo)	3-4	High	1,200 simulation runs	Moderate-High
BLADE-MARITIME, BLADE-INFRA	Design complete	3	High	Simulation	Moderate
BLADE-SPACE	Preliminary design	2-3	High (15-doc package)	Specified, not executed	Low-Moderate
BLADE-CUAS / AGENT-HSM / SWARM / INFRA-OT / FINANCE	Emerging platforms	2-4	Medium-High	Simulation / emulator	Low-Moderate
Rover and UAV testbeds	Design complete, build in progress	3-4	High	Simulation (350 / 250 runs)	Moderate
Software reference architectures (6)	Reference architecture	3-4	Medium-High	Specification + simulator	Low-Moderate
Simulation portfolio (19)	Operational research tools	Tooling	High	Reproducible, seeded runs	High (as tools)
Publications, patents, books	Published / disclosed	n/a	High	Public identifiers (DOI, SSRN, LCCN)	High (as disclosure)

EVIDENCE AND DOCUMENTATION MATRIX

Y indicates the artifact exists and is publicly retrievable. Partial indicates incomplete coverage. Planned or Not yet indicate items identified for development. References are public site paths.

ITEM	PAPER	SIM	CODE	DIAGRAMS	DATA	VALIDATION	REFERENCE
AUTHREX pipeline	Y (TSD)	Y	Y	Y	Y	Sim + formal	authrex.systems
SATA, HMAA, CARA, FLAME	Y (Zenodo)	Y	Y	Y	Y	Sim; HMAA TLA+	burakoktenli.com
MAIVA, ADARA	Y (Zenodo)	Y	Y	Y	Y	Simulation	burakoktenli.com
ERAM	Y (SSRN)	Y	Y	Y	Y	Sim (Monte Carlo)	burakoktenli.com/eram
BLADE-EDGE	Y (Zenodo)	Y	Y (BOM/ICD)	Y	Y	Sim / design	burakoktenli.com/blade-edge
BLADE-AV	Y (Zenodo)	Y	Y	Y	Y	1,200 runs	burakoktenli.com/blade-av
BLADE-MARITIME, INFRA	Y (Zenodo)	Y	Y	Y	Y	Simulation	burakoktenli.com
BLADE-SPACE	Y (15-doc)	Planned	Y (BOM)	Y	Not yet	Specified	burakoktenli.com/blade-space
Emerging platforms (5)	Y	Partial	Y	Y	Partial	Sim / emulator	burakoktenli.com
Software set (6)	Y (spec)	Y	Partial	Y	Y	Simulator	authrex.systems
Simulations (19)	Y (guides)	Y	Y	Y	Y	Reproducible	authrex.systems
Publications, books	Y	n/a	n/a	n/a	n/a	DOI / LCCN	burakoktenli.com

[EVIDENCE CLASSIFICATION]

EVIDENCE CLASS	WHAT BELONGS HERE
Supported	Public deposits and working papers with registered identifiers; reproducible seeded simulations; the HMAA TLA+ verification; bill-of-materials-level design packages.
Partial	Cross-domain portability shown in simulation; standards mappings as design targets; the simulated no-unsafe-action results, which hold in simulation only.
Conceptual	BLADE-SPACE at preliminary design; full hardware enforcement at the actuator boundary; any operational command-and-control relevance.
Unsupported until tested	Deployment readiness; formal qualification; government or agency approval; field or hardware performance. None of these is claimed in this catalog.

MISSION ALIGNMENT MATRIX

Alignment is marked only where it is genuinely supported. A filled circle marks a primary fit; a half circle marks a contributing or partial fit; a blank cell indicates no claimed alignment.

ITEM	AI	CYBER	INFRA	SPACE	DEF	RES	H - M	NOTES
AUTHREX pipeline	●	●	●	●	●	●	●	Domain-agnostic core
SATA, HMAA, CARA, FLAME	●	◐	◐	◐	●	◐	●	Core frameworks
MAIVA	●		◐	●	◐		◐	Swarm consensus
ADARA	●	●		◐	●		◐	Deception / EW resilience
ERAM	●			◐	●		◐	C2 escalation
BLADE-EDGE					●		◐	High-consequence defense
BLADE-AV	◐	◐				●	●	Automotive ASIL-D
BLADE-MARITIME	◐				●	◐	◐	Maritime surveillance
BLADE-INFRA	◐	◐	●		◐	●		ICS / SCADA
BLADE-INFRA-OT		◐	●			●	●	IT/OT boundary
BLADE-SPACE	◐	◐		●	◐	◐		Orbital autonomy
BLADE-CUAS	◐	◐			●	◐	●	Counter-UAS
BLADE-AGENT-HSM	●	●					◐	Hardware root of trust
BLADE-SWARM	●	◐		●	◐		◐	Swarm autonomy
BLADE-FINANCE	●	◐	◐			●		Economic security
AUTHREX-AGENT	●	●	◐				◐	Agentic AI
AUTHREX-AGENT-CYBER	●	●	◐		◐	◐		Cyber-defense (governance only)

Legend: AI = AI and autonomy; Cyber = cyber and security; Infra = critical infrastructure; Space = space and aerospace; Def = defense; Res = resilience; H-M = human-machine systems.

RISK AND LIMITATION ASSESSMENT

This section states the program limitations directly. It is written for a technical reviewer who needs an accurate picture of what has and has not been demonstrated.

[TECHNICAL LIMITATIONS]

The architecture is demonstrated in seeded simulation and, for the HMAA state machine, in formal verification. The central design premise, that authority enforcement terminates at a hardware boundary that software compromise cannot bypass, is specified but has not been demonstrated on physical silicon. No platform has been built. Performance, timing, and fault behavior on real hardware are therefore projected from design and simulation, not measured.

[VALIDATION GAPS]

There has been no hardware-in-the-loop testing and no independent red-team evaluation. Reported run counts and the zero-unsafe-actions results are simulation outcomes under seeded conditions, not field results. Formal verification depth is uneven: HMAA is verified across 23,748 distinct reachable states at depth 9, while other frameworks carry specifications or self-test suites rather than full state-space proofs. BLADE-SPACE has a specified but unexecuted verification campaign.

[DOCUMENTATION AND REVIEW GAPS]

The work is single-author. One component protocol, SATA, has been published in a peer-reviewed journal (Journal of Hardware and Systems Security, Springer Nature, 2026); the broader corpus of open-access deposits and SSRN working papers establishes public disclosure and priority but is not peer-reviewed. Independent third-party citation is currently limited. Standards mappings are research artifacts that describe how the frameworks relate to those standards; they are not accreditations, audits, or designated-engineering-representative findings.

[KEY ASSUMPTIONS]

Sensor-trust computation can be made fast and reliable enough to gate authority in real time on target hardware. The hardware authority boundary can be integrated at the actuator level on the intended FPGA or ASIC platforms. Simulation fault models are representative of the relevant real-world failure and adversarial conditions. Provisional patent applications can be converted to non-provisional filings within the statutory window.

[MISSING EVIDENCE]

Built hardware for any platform, and measured (not simulated) interlock behavior. An independent red-team report against sensor spoofing, authority hijack, Byzantine compromise, and jamming. Peer-reviewed publication of the core frameworks and the ERAM model. Any operational or third-party deployment evidence; BLADE-FINANCE uses synthetic data only.

PRIORITY IMPROVEMENT PLAN




Figure 6. Development roadmap from the current position (self-assessed TRL 2 to 4) to the near-term target (self-assessed TRL 6). The sequence runs from reference confirmation and language standardization, through testbed builds and the FPGA governance bitstream, to hardware-in-the-loop testing, independent red-team evaluation, and peer-reviewed submission.

[IMMEDIATE (BEFORE WIDER DISTRIBUTION)]

Standardize the HMAA authority-tier nomenclature across both public sites to the four-level model (A3 to A0 / T3 to T0). Confirm that every cited DOI, SSRN identifier, and patent receipt resolves to a live, correct record. Ensure no public claim states or implies field or hardware results where only simulation evidence exists. Reconcile the small count discrepancies between the two sites to a single canonical figure.

[HIGH PRIORITY (FOR SERIOUS TECHNICAL REVIEW)]

Form a small team to enable an independent red-team evaluation and independent replication of published results; these are the largest credibility gaps at present. The first peer-reviewed journal article was published in July 2026; pursue additional peer-reviewed submissions. Convert the seven provisional patent applications to non-provisional filings before the earliest window closes in approximately March 2027. Pursue SBIR Phase I funding to commission the FPGA governance bitstream and perform hardware-in-the-loop validation of the SATA-FLAME pipeline.

[MEDIUM PRIORITY (FOR A STRONGER PRESENTATION)]

Build the Rover and UAV testbeds, whose designs are complete, and report measured results alongside the existing simulation data. Execute the specified BLADE-SPACE verification and validation campaign. Submit the core frameworks and the ERAM model to peer-reviewed venues and track outcomes honestly. Expand the independent citation footprint through targeted dissemination.

[LONG-TERM (FOR FUTURE MATURITY)]

Advance multi-framework governance on physical hardware from self-assessed TRL 3-4 toward self-assessed TRL 6. Pursue formal qualification pathways where relevant (ISO 26262 ASIL-D for automotive, MIL-STD-882E for defense). Establish a research partnership, OEM adoption, or cooperative research agreement to validate in a relevant operational context. Extend the architecture to additional domains only where a genuine need and a credible validation path exist.

FINAL PROFESSIONAL SUMMARY AND ENGAGEMENT

[SUMMARY]

The work develops a single idea, authority as an engineered lifecycle, and carries it from formal specification through simulation to a family of domain-specific reference designs. The architecture is coherent across all of its instantiations, and the documentation is unusually complete for an independent research program: every system has a public deposit, a reproducible simulation, and a defined interface. The strongest areas are the coherence of the cross-domain architecture, the formal verification of the HMAA authority state machine, the breadth and reproducibility of the simulation portfolio, and the honesty of the readiness posture. The most important next steps convert simulation evidence into measured hardware evidence and add independent scrutiny; none of these change the architecture, they raise the level of proof behind it.

[HOW THE WORK CAN BE REVIEWED]

Every system in this catalog has a public artifact. The frameworks and platforms are deposited on Zenodo with registered DOIs; the ERAM model and related papers are on SSRN; the seven provisional patent applications are filed with the USPTO; and the simulations run client-side in a browser with a seeded pseudo-random number generator, so a reviewer can reproduce a result independently.

[FORMS OF ENGAGEMENT THAT FIT THE CURRENT MATURITY]

Technical review of a specific framework or platform against a defined assurance standard (for example, HMAA against a formal-methods checklist, or BLADE-AV against an ISO 26262 ASIL-D pathway). A cooperative research arrangement or sponsored evaluation to commission the FPGA governance bitstream and perform hardware-in-the-loop validation on the Rover or UAV testbed. An independent red-team evaluation against sensor spoofing, authority hijack, Byzantine compromise, and jamming. SBIR or other early-stage instruments aligned with a named subtopic. Co-authorship or peer review of the core frameworks and the ERAM model. Written technical exchange is the preferred starting point; a scoped review of a single system is a more useful first step than a broad survey of the whole portfolio.

[CORRESPONDENCE]

BURAK OKTENLI · INDEPENDENT RESEARCHER, AI GOVERNANCE AND SAFETY-CRITICAL AUTONOMOUS SYSTEMS
AUTHORITY & ARCHITECTURE IMPRINT · WASHINGTON, DC · INFO@BURAKOKTENLI.COM · ORCID 0009-0001-8573-1667
PROJECT SITES: BURAKOKTENLI.COM · AUTHREX.SYSTEMS · PUBLIC RECORDS: ZENODO · SSRN · USPTO PATENT CENTER

APPENDIX A: GLOSSARY

ADARA	Adversarial Deception-Aware Reasoning; adjusts authority by a computed deception probability.	HSM	Hardware Security Module; a tamper-resistant device that protects keys and authority state.
AISBOM	AI Software Bill of Materials; an inventory of an AI system's components emitted by the audit ledger.	ICS / SCADA	Industrial Control Systems and Supervisory Control and Data Acquisition.
ASIL-D	Automotive Safety Integrity Level D; the highest integrity level under ISO 26262.	ISA/IEC 62443	A standards series for industrial automation and control-system security.
BFT	Byzantine Fault Tolerance; consensus that tolerates arbitrarily faulty or malicious nodes.	IT/OT	Information Technology and Operational Technology; the corporate and control-system domains.
BLADE	The family of domain-specific hardware governance platforms in this catalog.	MAIVA	Multi-Agent Integrity Voting; Byzantine-resilient multi-agent consensus.
BOM	Bill of Materials; the itemized component list and cost for a hardware platform.	MIL-STD-810G / 882E	DoD standards for environmental engineering and for system safety.
CARA	Controlled Authority Recovery Architecture; deterministic recovery from authority lockout via GREP phases.	NDA	National Defense Authorization Act; annual U.S. defense policy legislation.
CC EAL6+	Common Criteria Evaluation Assurance Level 6 augmented; a high hardware-security assurance level.	NERC CIP	North American Electric Reliability Corporation Critical Infrastructure Protection standards.
CUAS	Counter-Unmanned Aircraft Systems; defense against small uncrewed aircraft.	NIST AI RMF	NIST Artificial Intelligence Risk Management Framework.
DoDD 3000.09	DoD Directive on autonomy in defense systems; requires appropriate human judgment over the use of force.	PBFT	Practical Byzantine Fault Tolerance; a classical BFT consensus protocol.
DO-178C / DO-333	RTCA standards for airborne software and its formal-methods supplement.	SATA	Sensor & Actor Trust Assessment; computes a continuous sensor-trust scalar.
Dempster-Shafer	An evidence-theory method for combining uncertain information from multiple sources.	SBIR / STTR	Small Business Innovation Research and Small Business Technology Transfer programs.
ECDSA P-256	Elliptic Curve Digital Signature Algorithm over the P-256 curve; used to sign audit records.	SLTT	State, Local, Tribal, and Territorial; non-federal government partners.
ERAM	Escalation Risk Assessment Model; quantifies escalation risk in AI-enabled command and control.	TLA+	A formal specification language used to verify the HMAA authority state machine.
FLAME	Forced Latency for Authority-Managed Escalation; imposes a deliberation window before irreversible action.	TPM	Trusted Platform Module; a hardware anchor for attestation and authority state.
FPGA	Field-Programmable Gate Array; reconfigurable hardware used for the governance boundary.	TRL	Technology Readiness Level; a nine-point scale for technology maturity.
GREP	Guard, Reduce, Evaluate, Promote; the phased recovery sequence used by CARA.	USV / UUV / UAS	Uncrewed Surface, Underwater, and Aircraft Systems.
HMAA	Human-Machine Authority Allocation; computes a graded authority level from trust and context.		

Records Added in Revisions 3.1 and 3.2

This addendum records developments since Version 3.0 and supersedes the corresponding counts and peer-review statements elsewhere in this document where they differ.

[PEER-REVIEWED JOURNAL ARTICLE]

Oktenli, B. SATA: Sensor Attestation and Trust Anchoring: A Formally Specified, Simulation-Evaluated Continuous-Trust Protocol for Autonomous-Systems Sensor Edges. Journal of Hardware and Systems Security (Springer Nature) 10, 17 (2026). DOI 10.1007/s41635-026-00190-4. Received 4 June 2026; accepted 14 July 2026; published 21 July 2026. This is the program's first peer-reviewed journal article; the published protocol corresponds to the SATA architecture.

[EXTERNAL EVALUATION RECORD · AIR FORCE RESEARCH LABORATORY]

A program white paper on conditions-based authority governance and assurance for the virtual command-and-control layer was favorably evaluated under Air Force Research Laboratory Broad Agency Announcement FA8750-24-S-7003 (CANVAS) and assessed as of interest to the Air Force; it was not funded at the time due to stated budget constraints (memorandum, June 2026). An evaluation record is not an endorsement, sponsorship, or contract.

[GOVERNOR CONSOLE REGISTER · THREE STANDALONE CONSOLES]

SIMULATION	GATE MODEL	METHOD	SIZE
AUTHREX-RTA Governor Console	Independent monitor · Simplex / ASTM F3269	Run-time-assurance invariant gate; SHA-256 hash-chained decision ledger; replay determinism; five stress scenarios; synthetic data	264 KB
AUTHREX-ENGAGE Governance Console	Fail-closed invariant monitor	Five adjudicated outcomes (authorize, clamp, downgrade, deny, hold for human); hash-chained ledger; five seeded scenarios; synthetic data	138 KB
AUTHREX-SAFING Safe-State Console	Fail-safe safing governor	Safe-state lifecycle (entry, hold, deepen, re-arm gated exit); safe-state floor never applied as an authority grant; deterministic replay against five published reference roots; synthetic data	354 KB

[SAFING V1.5.0 VERIFICATION RECORD]

The AUTHREX-SAFING console v1.5.0 carries a documented verification record: 1061 checks pass (S1 to S59); the formal authority model check enumerates 2,488,320 executions across 15 model states and 23 transitions to a maximum depth of 7 with all nine modeled properties holding; 11,739 terminal decisions map across three scenarios and fifty seeds with zero mismatches; 50 of 50 mandatory mutants killed; a 43-check adversarial battery passes with signed results. These are deterministic, seeded results on synthetic scenarios within the packaged build; they are not an independent validation, and no formal TRL determination has been performed for this console. Standalone sha256 ef8b828e039a12bdd3873723195bed87027df59d64ce82f91c72da09492e6bd6.

[PROTOTYPE CONSOLE RECORDS · ADDED IN REVISION 3.2]

AUTHREX-TEAM Team Authority Governance Console v0.16 (86 KB): a hardened interactive prototype, not a reproducible signed release; there is no release archive, no detached signature, and no signing key for this artifact. An independent arbiter over an untrusted allocator maintains exactly one eligible authenticated holder per responsibility, with a safe-hold fallback and an allowlisted supervisor override. Documented v0.16 evaluation: the boot self-test reports engine encapsulation; 20 of 20 scripted, seeded scenarios are contained with no observed unsafe authority transitions, zero authority gaps, and zero split-authority states; from seed 20260727 the build reproduces reference ledger root 5a937676527fbd33a9704978b772af925c18bfe56691dbabb659e23150a168c2. Standalone sha256 797076b55de768453aae13e74c677306eb4281a9911ed0875991d12ea01399e is the only integrity claim. Governance and assurance only; synthetic data; no formal TRL determination has been performed.

AUTHREX-REDLINE Positive Human Control Console v0.13 (117 KB): a hardened standalone prototype, not a reproducible signed release; there is no release archive, no detached signature, and no signing key for this artifact. The critical authorization can be held only by two distinct authenticated humans, a machine is never eligible, and the failure safe state is NO-GO; supervisor vetoes and separately attributed arbiter fail safe holds are recorded as what they are. Documented v0.13 evaluation: 27 of 27 bundled deterministic scenarios contained with all six audits passing; a seeded sweep of 250 seeds over 25,000 ticks records zero audit failures, zero machine key issuances, zero GO states without two distinct humans, and 843 fail safe holds, none attributed to a human; from seed 20260728 over 400 ticks the build reproduces reference ledger root 7bce87c43f67203852c3d02e3477cd75dec91f39e7f030fbff45d8f59a2fbb72. Standalone sha256 480459efb3cfc1023c6b50d843c98bfa9ed3d340f8a00eeb6308f0743914164 is the only integrity claim. Governance and assurance only; synthetic data; no formal TRL determination has been performed.

Scope unchanged. All three consoles and the TEAM and REDLINE prototypes are governance and assurance only: no vehicle dynamics, no actuation, no targeting, no kinetic effects, no live systems. All simulation counts in this document now read against a portfolio total of twenty-two catalogued browser simulations.

[DISCLAIMERS AND STATUS LABELS]

This catalog presents public-source, self-authored research artifacts. Evaluation is simulation-based unless explicitly stated otherwise. Hardware references are design specifications unless physical testing is specifically identified. The catalog does not represent official government endorsement, sponsorship, or affiliation.

Catalog status labels distinguish between operational physical hardware, browser-based simulations, formal models, reference designs, and future validation plans. A listed reference design should not be read as an approved product or fielded system.

End of document. This catalog is an independent research portfolio authored by Burak Oktenli under the Authority & Architecture imprint and is not a government publication. Best read alongside the program white paper and the per-system deposits, which contain the full bills of materials, simulation data, and source code.

AUTHREX RESEARCH PROGRAM

BURAK OKTENLI · PRINCIPAL RESEARCHER · WASHINGTON, DC
INFO@BURAKOKTENLI.COM · AUTHREX.SYSTEMS · BURAKOKTENLI.COM
ORCID 0009-0001-8573-1667

ERRATA AND CANONICAL RESTATEMENT

AUTHREX Technical Project Catalog V3.2 · appended August 2026 · formal coverage work package

This catalog states the 23,748-state result in five places without its scope clause. The figure is correct; the scope clause below applies to every occurrence.

Canonical HMAA result sentence, per AUTHREX-HYBRID-ASSURANCE v2.0.13 evidence state:

The authority state machine is specified in TLA+ and model checked with TLC across 23,748 distinct reachable states at depth 9 (26,397,356 states generated) with no counterexample; of 8 stated properties, 5 invariants and 1 liveness property held, and 2 properties are vacuous at this bound, one upgrade-path (UpgradelsStepwise) and one trust-authority consistency (TrustAuthorityWeakConsistency); the result describes the discrete authority automaton under the assumption that instantaneous authority equals its target and does not cover continuous behaviour between decision instants.

This page corrects the statements above without altering the original pages, which are retained verbatim as issued. Nothing in this document claims hardware validation, certification, or verification by an independent party; every formal result is bounded and model-relative. Full status: authrex.systems/formal-coverage.html

ERRATA 2: WITHDRAWN SIMULATION-RUN CLAIMS

AUTHREX Technical Project Catalog V3.2 - appended August 2026

The following claims appearing in this catalog are WITHDRAWN and must not be cited:

1. BLADE-AV: "1,200 simulation runs, no observed unsafe authority transitions" and all variants of that sentence.
2. Rover Testbed: "350 simulation runs, 7 fault scenarios, no observed unsafe authority transitions".
3. HMAA-UAV Platform: "250 simulation runs".
4. The aggregate figure "2,800+ simulation runs, no observed unsafe authority transitions" across flight-relevant platforms, which sums the three counts above.
5. Every paired assertion of the form "no observed unsafe authority transitions" or "no unsafe actions recorded".

Reason for withdrawal.

The underlying run record is not reproducible from the published artifacts. A run count and a paired safety assertion together read as an empirical result; without a reproducible record, no such result is supported. The counts and the assertion are withdrawn together because the assertion derives its weight from the count.

What replaces them.

Scenario counts, which are properties of the packaged configuration and are checkable by inspection: BLADE-AV twelve attack scenarios, Rover Testbed seven fault scenarios, HMAA-UAV five adversarial scenarios, each demonstrated in simulation. Component counts, bill-of-materials figures, DOIs, and TLA+ state figures in this catalog are unaffected and stand as published. The seeded pseudo-random generator statement is a methodology statement, not a result, and also stands.

This page corrects the pages above without altering them; the original text is retained verbatim as issued. Nothing in this document claims hardware validation, certification, or verification by an independent party.