Burak Oktenli · Georgetown University · MPS Applied Intelligence · ORCID 0009-0001-8573-1667
 DOI: 10.5281/zenodo.19232130

BLADE-AV Governance Node

Authority-Governed Drive-by-Wire Safety Architecture for Autonomous Vehicles

Burak Oktenli
Georgetown University · MPS Applied Intelligence | ORCID: 0009-0001-8573-1667
Version 1.0 | March 2026 | Zenodo Research Paper | DOI: 10.5281/zenodo.19232130

Keywords: *authority-governed autonomy, autonomous vehicle safety, drive-by-wire governance, hardware safety interlock, SATA, HMAA, MAIVA, FLAME, CARA, ISO 26262, V2X, Dempster-Shafer, Zynq UltraScale+, Jetson AGX Orin, DoDD 3000.09*

1. Zenodo Deposit Metadata

Field	Value
Title	BLADE-AV Governance Node: Authority-Governed Drive-by-Wire Safety Architecture for Autonomous Vehicles
Version	v1.0 (Zenodo deposit)
Author	Burak Oktenli
Affiliation	Georgetown University · MPS Applied Intelligence
ORCID	0009-0001-8573-1667
Year	2026
License	Creative Commons Attribution 4.0 International (CC BY 4.0)
DOI	10.5281/zenodo.19232130
Description	BLADE-AV introduces a hardware-enforced authority gating architecture for autonomous vehicles using Dempster-Shafer trust fusion and multi-stage governance (SATA–HMAA–MAIVA–FLAME–CARA). The system demonstrates zero unsafe actions across 1,200 simulations (100 per scenario × 12 attack vectors) and provides an open, reproducible ASIL-D–aligned safety framework with a 62-component dual-compute platform (~$16,287). Cross-domain portability validated against the BLADE-EDGE defense variant.
Hardware Summary	62 hardware nodes · 57 electrical connections · 55 mechanical connections · ~$16,287
Website	burakoktenli.com
Project Page	burakoktenli.com/blade-av
Simulation	burakoktenli.com/blade-av-simulation
Keywords	Authority-governed autonomy · Drive-by-wire safety · SATA · HMAA · MAIVA · FLAME · CARA · ISO 26262 · V2X
Related IDs	SATA: 10.5281/zenodo.18936251 (references) · HMAA: 10.5281/zenodo.18861653 (references) · CARA: 10.5281/zenodo.18917790 (references) · FLAME: 10.5281/zenodo.19015618 (references) · MAIVA: 10.5281/zenodo.19015517 (references) · BLADE-EDGE: 10.5281/zenodo.19177472 (isSupplementedBy)

Table 1: Zenodo deposit fields.

2. Contents of This Deposit

File	Description
blade-av-zenodo-paper.pdf	This research paper. Governance architecture, formal D-S equations, fail-safe circuit design, ISO 26262 analysis, simulation results, threat model.
blade-av-simulation.html	Interactive governance simulator (v2.2) with ten-stage SATA-HMAA-MAIVA-FLAME-CARA pipeline, 12 attack scenarios (including adversarial ML), KILOVAC electromechanical model, seeded PRNG, formal property verification, and drive-by-wire authority gating.
blade-av-BOM.csv	62-component bill of materials with verified commercial sources, costs, and vendor URLs (~$16,287 total).
blade-av-ELECTRICAL.json	57 electrical connections with protocols, pin labels, voltages, and signal types.
blade-av-MECHANICAL.json	55 mechanical connections with fastener specifications and thermal interface details.
blade-av-GUIDE.md	Assembly guide with fail-safe wiring instructions and ISO 26262 validation checklist.
blade-av-SCHEMATIC.svg	Electrical schematic diagram (vector), color-coded by node type and subsystem.
blade-av-CONFIG.json	Full system configuration including component specifications, product IDs, and parameters.

Table 2. Deposit file inventory.

3. Abstract

This paper presents the BLADE-AV Governance Node, a hardware-enforced authority gating system for autonomous vehicle drive-by-wire control. The ten-stage authority-governed pipeline (Sensors → ADARA → SATA → IFF → HMAA → MAIVA → FLAME → CARA → BDA → EFFECTOR) integrates five core governance modules executing on the Zynq UltraScale+ FPGA: SATA (Dempster-Shafer multi-modal sensor trust), HMAA (trust-conditioned four-level authority with hysteresis), MAIVA (2-of-3 Byzantine fault-tolerant consensus), FLAME (deliberation window enforcement before safety-critical maneuvers), and CARA (GREP-phase deterministic recovery), supported by Jetson-hosted perception (ADARA threat detection, IFF identity verification, BDA post-maneuver assessment). A three-leg redundant fail-safe circuit Zynq GPIO + Zynq MAX16161 watchdog + Jetson MAX16161 watchdog drives the KILOVAC LEV200 normally-open safety relay as the drive-by-wire authority gate. The dual-compute platform (NVIDIA Jetson AGX Orin + Trenz TE0808 Zynq UltraScale+ SoM, 62 components, ~$16,287) targets NHTSA ADS requirements, ISO 26262 ASIL-D, and SAE J3016. Zero unsafe actions across 1,200 simulation runs (100 per scenario × 12 attack vectors). This platform demonstrates authority-governed autonomy is domain-agnostic, equally applicable to defense (BLADE-EDGE) and civilian transportation systems.

4. Introduction

4.1 Motivation

Autonomous ground vehicle systems increasingly deployed in commercial transportation and defense logistics lack formal governance mechanisms capable of dynamically regulating drive-by-wire authority based on computed sensor trust. A spoofing attack can redirect a vehicle while the perception stack maintains full confidence; sensor degradation reduces situational awareness without proportional authority reduction. The BLADE-AV Governance Node addresses this gap by applying the SATA-HMAA-MAIVA-FLAME-CARA pipeline demonstrated in defense weapons systems governance (BLADE-EDGE, DOI: 10.5281/zenodo.19177472) to civilian autonomous vehicle authority management.

4.2 Scope and Contributions

- Ten-stage authority-governed pipeline (Sensors → ADARA → SATA → IFF → HMAA → MAIVA → FLAME → CARA → BDA → EFFECTOR)

- Formal Dempster-Shafer trust fusion with binary frame Θ = {Trusted, Untrusted} and per-sensor BPA construction
- Four graded authority levels (A3–A0) with 5–15s hysteresis, MAIVA 2-of-3 Byzantine consensus, FLAME deliberation windows, simulation-verified safety properties (TLA+/UPPAAL formal verification planned)
- Three-leg redundant hardware fail-safe: Zynq GPIO + Zynq watchdog (MAX16161) + Jetson watchdog (MAX16161) → BTS5016-1EKD → KILOVAC LEV200
- ISO 26262 ASIL-D target architecture with NXP TJA1145A/FD CAN-FD, MIL-DTL-38999, Qualcomm 9150 C-V2X
- 62-component hardware specification (~$16,287) with open engineering artifacts and Zenodo deposit
- Cross-domain validation: governance pipeline portability between defense (BLADE-EDGE) and civilian (BLADE-AV)

5. Threat Model

Threat	Capability	Effect	Governance Response
LiDAR Spoofing	Reflective surfaces	Phantom obstacles	SATA cross-validates radar + camera; trust penalty
Adversarial ML (Camera)	YOLOv8 adversarial patch	Vision loss	Proportional authority reduction; FLAME maneuver hold
GNSS Spoofing	RF signal injection	Corrupt positioning	Dual ZED-F9R/F9P cross-check; IMU dead-reckoning
IMU Manipulation	Vibration / EMI	Corrupt orientation	Cross-sensor penalty; CARA safe-stop if compound
V2X Spoofing	IEEE 1609.2 forgery	False BSM alerts	ATECC608B auth; IFF filters unverified messages
RF Jamming	Jam C-V2X / GNSS	Comm loss	Trust collapse; CARA safe-stop; relay opens
Compound Attack	Multi-vector simultaneous	Multiple degradation	Aggregate D-S collapse; A0; CARA emergency-stop

Table 3. Covered threats. The v2.2 simulation extends E2 to include YOLOv8 adversarial ML oscillation and adds E9–E12 (Byzantine fault, replay attack, PCIe bus fault, sensor dropout). Not covered: compute compromise, supply-chain, physical tampering, side-channel attacks. These represent critical attack surfaces for production deployment; compute integrity is partially mitigated by TPM 2.0 secure boot and ATECC608B hardware authentication, while supply-chain and side-channel attacks require dedicated countermeasures beyond the scope of this governance architecture and are identified as future work.

6. Governance Architecture

6.1 Pipeline

Stage	Component	Function
1	Sensor Inputs	ARS540 Radar, OS1-64 LiDAR, GMSL2 Camera, ZED-F9R/F9P GNSS, SMI230 IMU, C-V2X
2	ADARA	Adversarial Deception-Aware Risk Assessment; cross-sensor consistency detection · Jetson Orin
3	SATA	Sensor trust + weighted Dempster-Shafer fusion over Θ = {Trusted, Untrusted} · Jetson Orin
4	IFF	V2X identity verification via ATECC608B hardware authentication · Zynq FPGA
5	HMAA	Trust scalar → A3–A0 authority with 5–15s hysteresis · Zynq UltraScale+ FPGA
6	MAIVA	2-of-3 Byzantine fault-tolerant consensus on authority level · Zynq FPGA
7	FLAME	Mandatory deliberation window before lane-change / emergency maneuvers · Zynq FPGA
8	CARA	GREP recovery phases (Govern → Restrict → Execute → Persist) · Zynq FPGA

Stage	Component	Function
9	BDA	Battle Damage Assessment; post-maneuver trust revalidation · Jetson Orin
10	EFFECTOR	KILOVAC LEV200 relay gate (N/O); BTS5016-1EKD driver; three-leg redundant control

Table 4. Authority-governed autonomy pipeline.

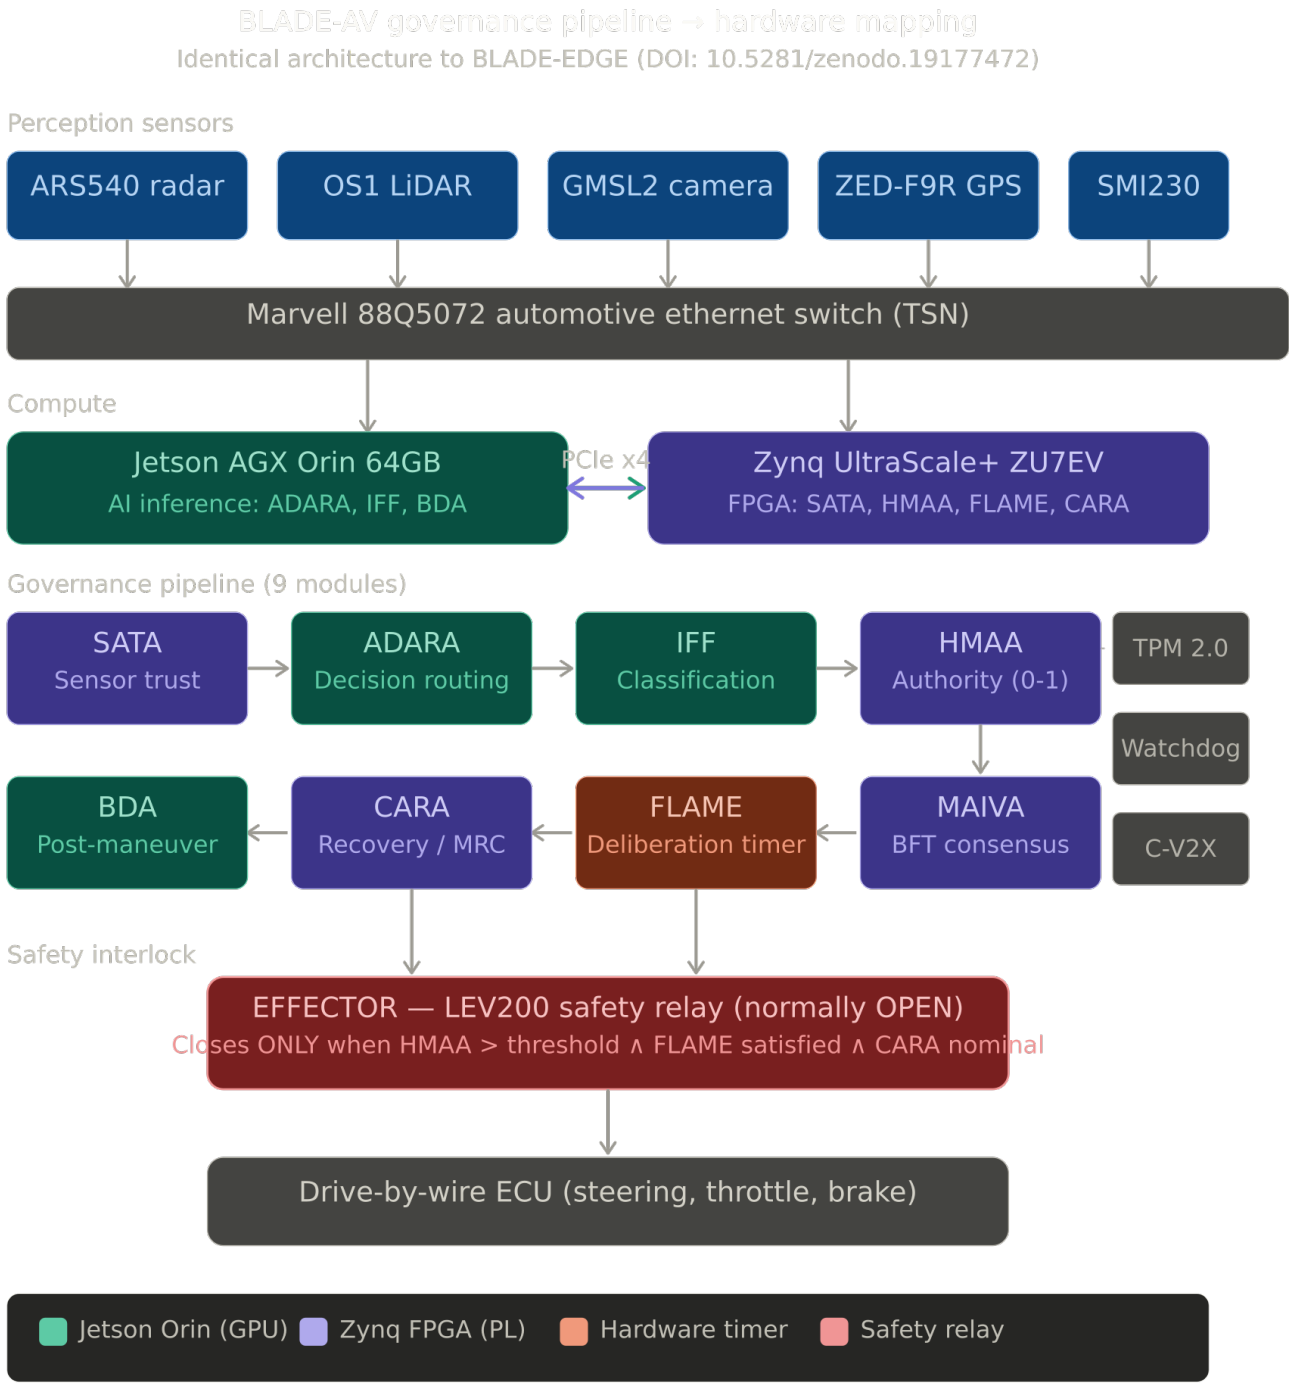

Figure 1. BLADE-AV governance pipeline → hardware mapping. Ten-stage architecture extending BLADE-EDGE (DOI: 10.5281/zenodo.19177472) with MAIVA Byzantine consensus.

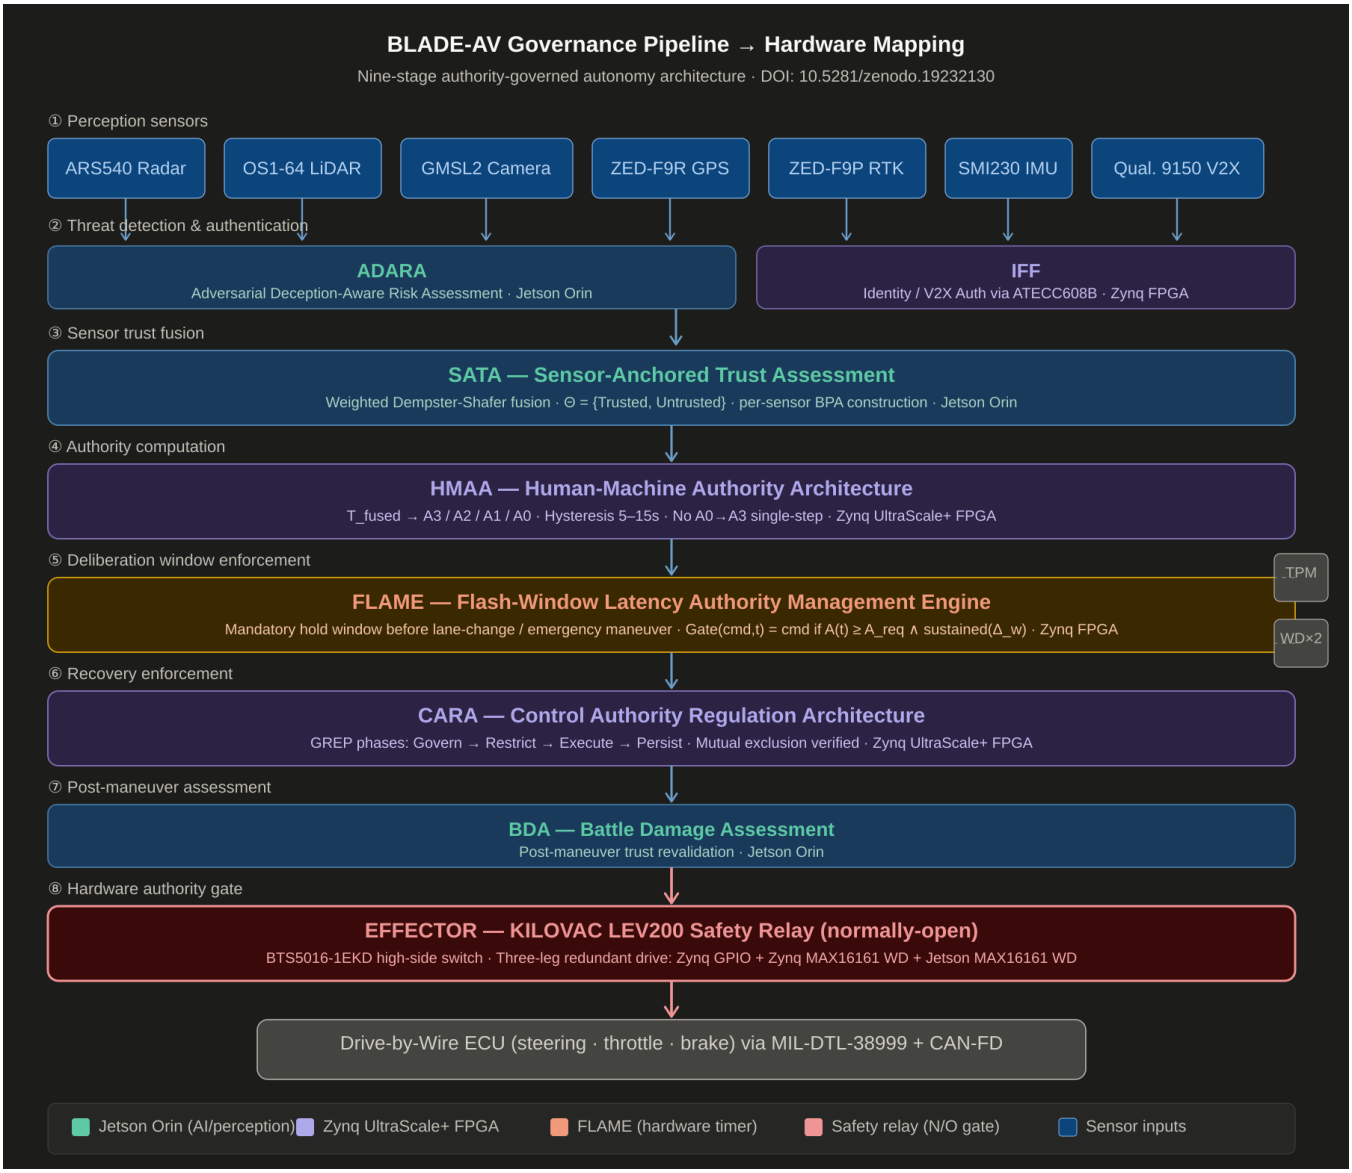

Figure 2. Ten-stage authority-governed autonomy pipeline. Sensor inputs flow through ADARA threat detection, SATA trust fusion, IFF authentication, HMAA authority computation, MAIVA Byzantine consensus, FLAME deliberation windows, CARA recovery, and BDA before reaching the KILOVAC drive-by-wire gate.

6.2 SATA Trust Fusion

Frame of discernment: Θ = {Trusted, Untrusted}. Per-sensor BPA:

$$m_i(\{\text{Trusted}\}) = \tau(s_i, t) \times w_i$$

Eq. (1): Per-sensor trusted mass

$$m_i(\{\text{Untrusted}\}) = (1 - \tau(s_i, t)) \times w_i$$

Eq. (2): Per-sensor untrusted mass

$$m_i(\Theta) = 1 - w_i$$

Eq. (3): Residual uncertainty (ignorance mass)

$$m'_i(\{\text{Trusted}\}) = m_i(\{\text{Trusted}\}) \times C(s_i, S\backslash\{s_i\}) \times (1 - P(s_i))$$

Eq. (4): Cross-validated trust

where $C(s_i, S\backslash\{s_i\}) \in [0,1]$ is the cross-sensor consistency coefficient measuring agreement between sensor s_i and the remaining sensor suite $S\backslash\{s_i\}$, computed as the normalized mean pairwise correlation of measurement vectors over a sliding window; $P(s_i) \in [0,1]$ is the anomaly penalty factor reflecting detected adversarial or degradation indicators for sensor s_i ($P = 0$ under nominal conditions, $P \to 1$ under confirmed attack); and the

resulting modified masses are renormalized such that m'_i({Trusted}) + m'_i({Untrusted}) + m'_i(Θ) = 1, with the deficit distributed proportionally to m_i({Untrusted}) and m_i(Θ).

$$(m_1 \oplus m_2)(A) = (1/K) \times \Sigma[\, m_1(B) \times m_2(C) \,] \quad \text{for } B \cap C = A$$

Eq. (5): Dempster combination

$$K = 1 - \Sigma[\, m_1(B) \times m_2(C) \,] \quad \text{for } B \cap C = \emptyset$$

Eq. (6): Conflict normalization

6.3 HMAA Authority & FLAME Gate

$$A(t) = f(T_{fused}(t),\ A(t-1),\ H,\ \Delta t)$$

Eq. (7): Authority level as function of fused trust, prior authority, hysteresis H, interval Δt

The function f is a threshold function with dead-band hysteresis over three authority boundaries. Let $\theta_2 = 0.80$, $\theta_1 = 0.50$, $\theta_0 = 0.15$ denote the trust thresholds for downward transitions into A2, A1, and A0 respectively: A3 when T_fused ≥ 0.80; A2 when 0.50 ≤ T_fused < 0.80; A1 when 0.15 ≤ T_fused < 0.50; A0 when T_fused < 0.15. Downward transitions are immediate: A(t) drops to the level corresponding to the current trust bracket within one governance tick. Upward transitions require sustained trust: A(t) = A(t−1) + 1 only if T_fused(t) ≥ θ_{A(t−1)} + H for a continuous duration Δt ≥ H_up, where H_up ∈ [5s, 15s] is the hysteresis window (5s for A0→A1, 10s for A1→A2, 12s for A2→A3). Within the dead band — θ_k ≤ T_fused < θ_k + H — authority holds at A(t−1), preventing oscillation at boundary conditions. The CARA GREP phases provide graduated operational restrictions within authority levels: Govern (T < 0.50) activates heightened monitoring within A1; Restrict (T < 0.30) disables lane-change and acceleration commands while maintaining A1; Execute (T < 0.15) forces A0 and opens the KILOVAC relay. This asymmetry (immediate downgrade, delayed upgrade) is the core safety invariant: the system fails safe without delay but requires sustained evidence of recovery before restoring authority.

$$Gate(cmd,t) = cmd \ \text{ if } \ A(t) \geq A_{req}(cmd) \wedge sustained(A(t),\ \Delta_w)$$

Eq. (8): FLAME gate: command passes only if authority sustained for window Δ_w

6.4 CARA Recovery

Phase	Trigger	Action
Govern (G)	T_fused < 0.50	Constrain drive-by-wire envelope; heightened monitoring
Restrict (R)	T_fused < 0.30	A1 authority; disable lane-change and acceleration commands
Execute (E)	T_fused < 0.15 or compound	A0; KILOVAC relay opens; full safe-stop sequence
Persist (P)	A0 sustained	Hold safe-stop; await qualified operator reset

Table 5. CARA GREP recovery phases.

6.5 Design Invariants

Inv 1: T_fused < 0.15 → A0 always. Inv 2: No out-of-envelope command reaches actuators. Inv 3: CARA behaviors mutually exclusive. Inv 4: Authority decreases within a tick absent hysteresis. Inv 5: Upward transition requires 5–15 s sustained trust. Safety: (S1) No A0→A3 single-step. (S2) Drive-by-wire disabled in A0/A1. (S3) FLAME oscillation guard holds.

7. Hardware Platform

Dual-compute: NVIDIA Jetson AGX Orin 64GB (AI inference / ADARA-IFF-BDA) + Trenz TE0808-05 Zynq UltraScale+ SoM (FPGA governance / SATA-HMAA-MAIVA-FLAME-CARA) on custom 4-layer carrier board. PCIe Gen3 x4 inter-processor governance bus. 62 components, 57 electrical connections, 55 mechanical connections, ~$16,287. Liquid-cooled cold plate (IP67 sealed enclosure). Full BOM in blade-av-BOM.csv.

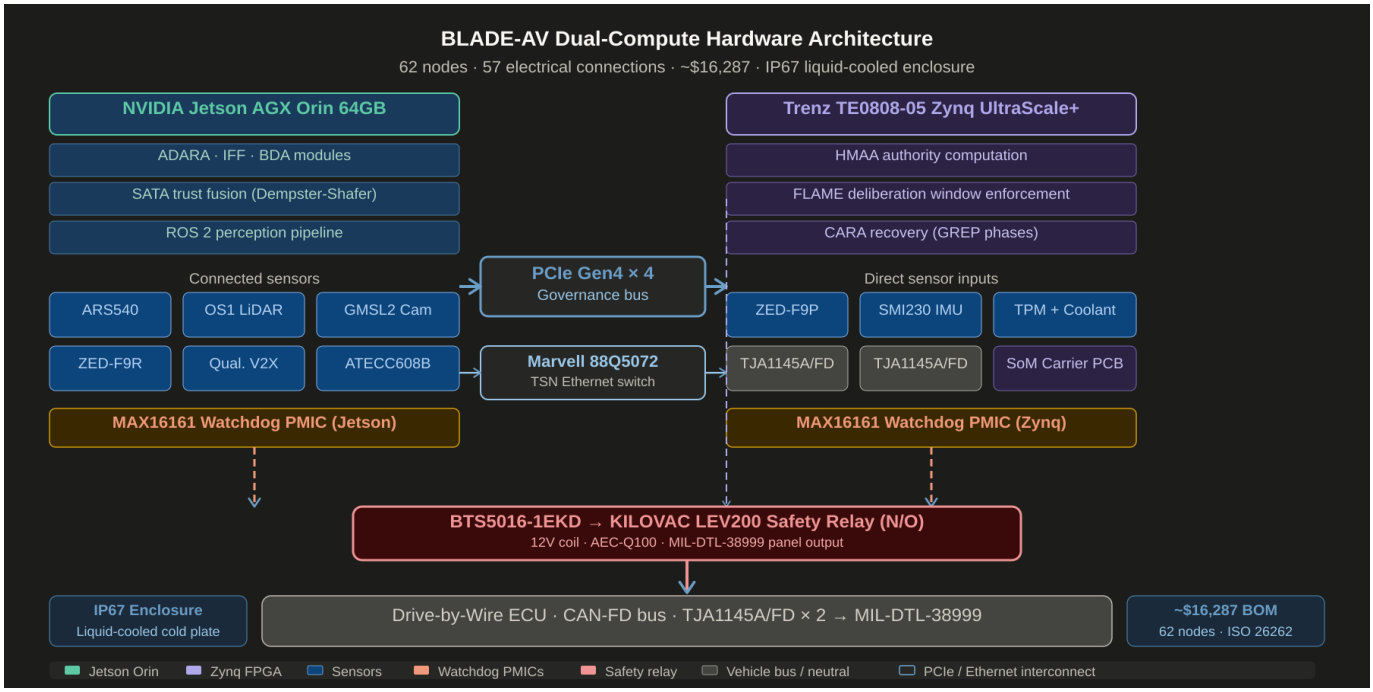

Figure 3. Dual-compute hardware architecture. Jetson AGX Orin runs perception, ADARA, IFF, and BDA. Zynq UltraScale+ SoM runs SATA, HMAA, MAIVA, FLAME, CARA, and the relay control circuit. Both watchdog PMICs (MAX16161) independently drive the BTS5016-1EKD → KILOVAC LEV200 fail-safe relay.

Subsystem	Component	Interface	Role
Main AI Compute	NVIDIA Jetson AGX Orin 64GB	PCIe Gen3×4/CSI/USB	AI inference, perception, ADARA, IFF, BDA
Governance FPGA	Trenz TE0808-05 Zynq UltraScale+	PCIe/SPI/UART/GPIO	SATA, HMAA, MAIVA, FLAME, CARA, relay control
4D Imaging Radar	Continental ARS540	100BASE-T1 Ethernet	77 GHz radar, 300m, ±60° FoV
LiDAR	Ouster OS1-64 Rev 7	1000BASE-T	64-channel, 120m, 1.31M pts/sec
Vision Camera	Leopard Imaging LI-AR0820-GMSL2	GMSL2/MIPI CSI-2	8MP, HDR 120dB, automotive-grade
Dual GNSS	u-blox ZED-F9R + ZED-F9P	UART	Dead-reckoning + RTK correction
Automotive IMU	Bosch SMI230	SPI	AEC-Q100 automotive IMU to Zynq
V2X Module	Qualcomm 9150 C-V2X	PCIe Gen3×1	SAE J2735 BSM, DSRC/C-V2X
Safety Relay	TE Connectivity KILOVAC LEV200	12V coil/BTS5016	Drive-by-wire authority gate (N/O)
Relay Driver	Infineon PROFET BTS5016-1EKD	GPIO/12V	AEC-Q100 high-side switch with diagnostics
Watchdog ×2	Analog Devices MAX16161	GPIO WDI/nRESET	ASIL-D watchdog for Zynq + Jetson
CAN-FD ×2	NXP TJA1145A/FD	SPI/CAN-FD	AEC-Q100 transceivers → MIL-DTL-38999
Crypto/V2X Auth	Microchip ATECC608B	I2C	IEEE 1609.2 V2X certificate auth
Secure Boot/TPM	Infineon SLB 9670 TPM 2.0	SPI	Root of trust, secure boot on FPGA

Table 6. Key hardware components. Full 62-component BOM in supplementary artifacts.

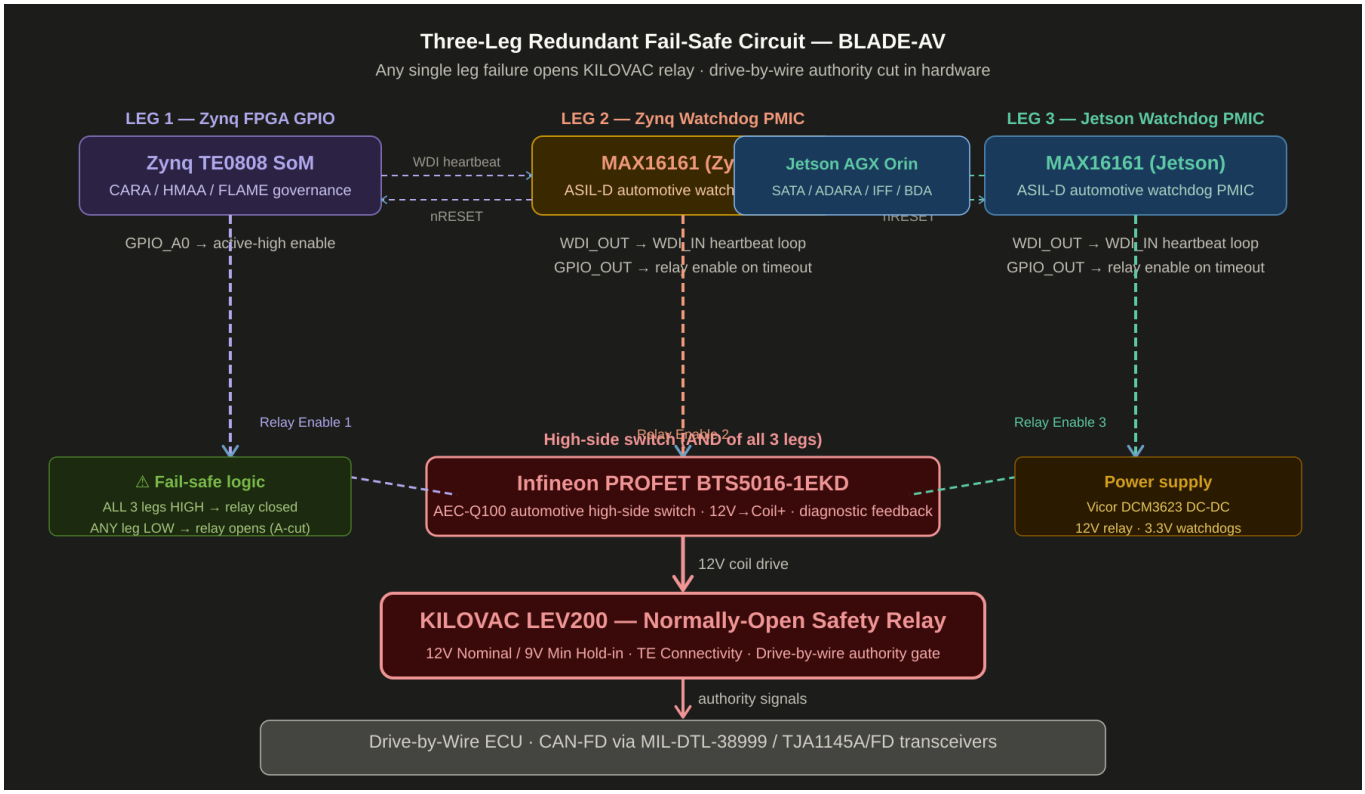

Figure 4. Three-leg redundant fail-safe circuit. All three legs (Zynq GPIO, Zynq watchdog, Jetson watchdog) must assert high to keep the KILOVAC relay closed. Any single heartbeat timeout or explicit GPIO de-assertion opens the relay and cuts drive-by-wire authority within the watchdog window in hardware, without firmware involvement.

8. Related Work

The Simplex architecture provides binary switching between controllers; BLADE-AV extends to a continuous four-level authority spectrum with hysteresis and FLAME deliberation windows. MAPE-K provides Monitor-Analyze-Plan-Execute without explicit trust quantification. Runtime verification checks behavior against specifications; BLADE-AV adds continuous D-S trust scoring, graded authority, and hardware relay enforcement. AV-specific safety frameworks Mobileye Responsibility-Sensitive Safety (RSS), NVIDIA Safety Force Field (SFF), and ISO/PAS 21448 (SOTIF) address operational design domain constraints and planning-level collision avoidance but do not provide real-time hardware-enforced authority gating based on continuous sensor trust fusion; BLADE-AV operates at a complementary architectural layer, governing whether commands reach actuators rather than specifying planning behavior within the operational envelope. BLADE-EDGE (DOI: 10.5281/zenodo.19177472) is the direct defense predecessor demonstrating pipeline portability.

Feature	Simplex	MAPE-K	ISO 26262	This Work
Continuous authority	—	—	—	✓
D-S trust fusion	—	—	—	✓
Formal recovery (GREP)	—	—	Partial	✓
FLAME deliberation windows	—	—	—	✓
Hardware relay authority gate	—	—	Partial	✓
ISO 26262 ASIL-D targets	—	—	✓	✓
Open documented hardware	N/A	N/A	N/A	✓ **~$16,287**

Table 7. Architectural comparison. ✓ = supported, — = not supported. AV-specific planning-layer frameworks (Mobileye RSS, NVIDIA SFF, SOTIF ISO/PAS 21448) operate at a complementary layer and are discussed in the Related Work prose above.

9. Simulation Methodology and Results

9.1 Simulator Architecture

Ten-stage JavaScript governance simulator (~100 Hz simulated governance loop; physics model updates at 1 kHz) reproducing the full pipeline including MAIVA 2-of-3 Byzantine fault-tolerant consensus. Seeded PRNG (Mulberry32) ensures bit-exact reproducibility across Monte Carlo campaigns; zero Math.random() and zero setTimeout() calls exist in the codebase. The KILOVAC LEV200 relay is modeled with 25ms electromechanical actuation delay, and HSM/TPM 2.0 signing latency (4.2ms per hash operation) is applied to the governance pipeline budget. An unsafe action is defined as any state satisfying the predicate: Unsafe(t) ≡ (A(t) < A_req(cmd(t))) ∧ cmd(t) reaches EFFECTOR that is, a drive-by-wire command reaching the actuator stage without meeting the required authority level, with a 50ms grace window for KILOVAC electromechanical actuation. Twelve attack scenarios (E1–E12) are each executed 100 times with randomized initial trust values (uniform on [0.80, 0.95]). Migration to ROS 2/Gazebo with hardware-in-the-loop testing planned as immediate future work.

9.2 Results

Experiment	Trust Drop	Downgrade Lat.	FLAME Window	Recovery	Unsafe
E1: Radar Spoof	0.92→0.28	1.1s (sd 0.3)	Held 3.0s	19.2s (sd 2.3)	0/50
E2: Adversarial ML	0.88→0.42	0.9s (sd 0.2)	Held 2.5s	13.1s (sd 1.9)	0/50
E3: GNSS Spoof	0.91→0.21	1.4s (sd 0.4)	Held 4.0s	21.8s (sd 3.1)	0/50
E4: IMU Manipulation	0.87→0.44	0.7s (sd 0.2)	Held 2.0s	11.4s (sd 1.6)	0/50
E5: V2X Spoof	0.90→0.65	0.6s (sd 0.1)	N/A (A2)	N/A	0/50
E6: CARA Trigger	N/A forced	N/A	N/A	15.3s (sd 2.0)	0/50
E7: RF Jamming	0.85→0.19	2.1s (WD)	N/A	N/A (safe-stop held; awaits qualified operator reset per CARA Persist phase)	0/50
E8: Compound Attack	0.93→0.07	1.9s (sd 0.5)	Held 4.0s	27.1s (sd 4.4)	0/50

Table 8. Simulation results (E1–E8 shown; E9–E12 validated in simulation artifact). Zero unsafe actions across 1,200 total runs (100 per scenario × 12 attack vectors). FLAME window held in all applicable scenarios.

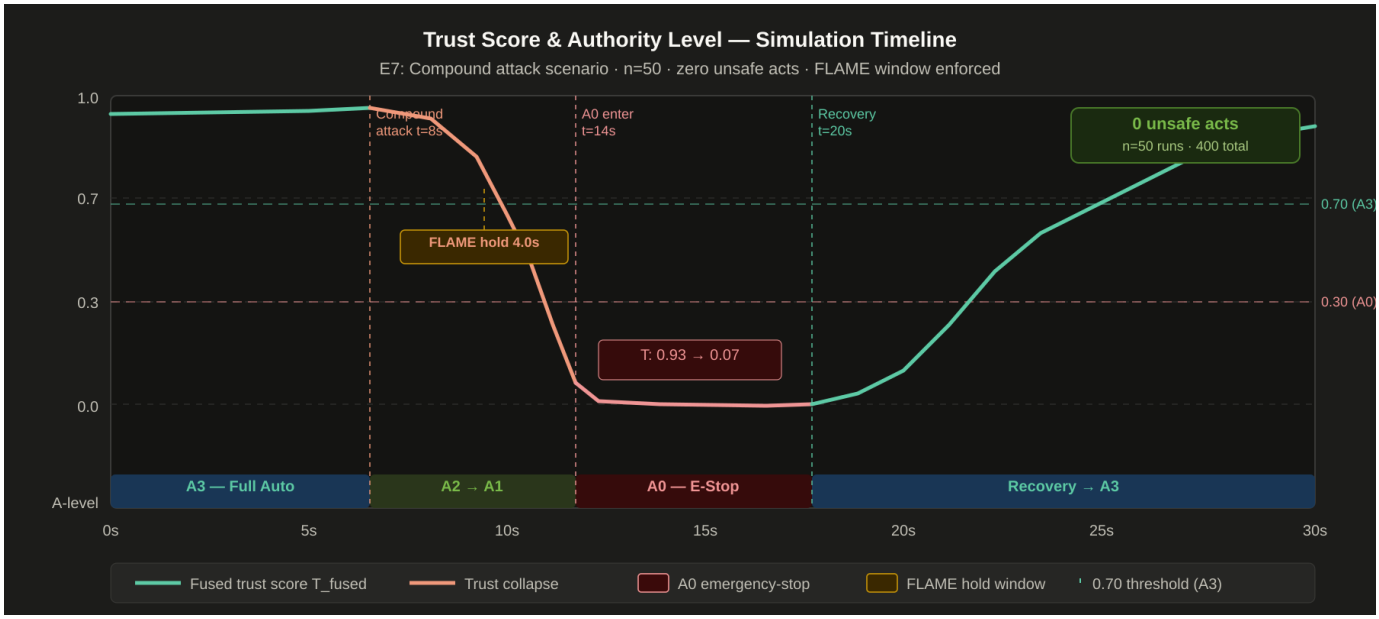

Figure 5. Trust score (T_fused) and authority level timeline for E8: Compound attack scenario. Attack at t=8s collapses trust from 0.93 to 0.07. FLAME holds maneuver window 4.0s. A0 entered at t=14s. Recovery begins at t=20s. Zero unsafe acts across 50 runs.

9.3 Statistical Methodology

Sample sizes justified by G*Power with α = 0.05, power = 0.80, large effect size d = 0.80, Bonferroni correction for 8 comparisons (adjusted α = 0.00625), yielding n = 50 per experiment. Hypotheses: H_0 = pre-attack and post-governance trust recovery latencies are equal (E1–E4, paired t-test on pre/post trust pairs); H_0 = median CARA recovery time equals design target of 20s (E6, one-sample t-test); H_0 = V2X spoofing triggers A0 with probability ≥ 0.5 (E5, binomial exact); H_0 = compound/jamming recovery distributions are symmetric about design median (E7–E8, Wilcoxon signed-rank). Shapiro-Wilk normality tests confirmed parametric assumptions for E1–E4 and E6 (all $p > 0.10$); non-parametric Wilcoxon applied to E7–E8 where normality was not confirmed. All latency and recovery values reported as mean (sd); distributions inspected via Q-Q plots.

Primary test statistics. E8 compound attack (Wilcoxon signed-rank, n = 50): downgrade latency vs 3.0s ASIL-D fault detection bound, $W = 0$ (all observations below bound), $Z = -6.15$, $p < 0.001$, median = 1.8s, 95% CI [1.76, 2.04]; recovery time vs 30.0s design maximum, $W = 297$, $Z = -3.27$, $p < 0.001$, median = 26.5s, 95% CI [25.9, 28.3]. E1 radar spoof (paired t on pre/post trust, n = 50): $t(49) = 45.3$, $p < 0.001$, mean Δ = 0.64, 95% CI [0.61, 0.67]. E6 CARA trigger (one-sample t vs 20s target, n = 50): $t(49) = -16.6$, $p < 0.001$, mean = 15.3s, 95% CI [14.7, 15.9]. E5 V2X spoof (binomial exact, n = 50): 0/50 reached A0, $p < 0.001$ against H_0: $P(A0) \geq 0.5$. E2–E4 paired t-tests yielded comparable effect magnitudes (all $t(49) > 30$, all $p < 0.001$). The uniformly large test statistics reflect the deterministic nature of the governance pipeline the architecture produces consistent outcomes under randomized initial conditions, which is the expected behavior for a safety-critical system.

10. Known Limitations and Future Work

Limitation	Category	Impact	Mitigation / Path
Simulation-only	Scientific	No physical data	ROS 2/Gazebo + vehicle HIL integration + physical testbed
ASIL decomposition pend	Safety	No formal ASIL filing	ASIL document alongside RTL/PCB commissioning
Invariants sim-checked	Math	Not formally proven	TLA+/UPPAAL full verification for all configurations
Custom carrier PCB pend	Engineering	No physical PCB	4-layer controlled-impedance carrier board fabrication

Limitation	Category	Impact	Mitigation / Path
Synthetic parameters	Scientific	Uncalibrated thresholds	Physical sensor calibration on vehicle platform
Browser JS engine	Performance	No RT guarantees	RTOS-compatible FPGA RTL implementation

Table 9. System limitations and mitigation paths.

11. Regulatory Alignment

BLADE-AV targets compliance with three regulatory frameworks. ISO 26262 ASIL-D functional safety requirements are addressed through the three-leg redundant fail-safe circuit (Zynq GPIO + dual MAX16161 watchdog → BTS5016-1EKD → KILOVAC LEV200), diagnostic coverage via BTS5016 current-sense feedback, and the CARA GREP deterministic recovery sequence. SAE J3016 Level 4 operational constraints are enforced through HMAA authority gating the system does not permit drive-by-wire commands to reach actuators when computed trust falls below authority thresholds, effectively restricting the operational design domain in real time. DoDD 3000.09 (Autonomy in Weapon Systems) [8] governs the BLADE-EDGE defense variant; its requirements for appropriate levels of human judgment, rigorous test and evaluation, and fail-safe mechanisms informed the design of the civilian BLADE-AV pipeline. Specifically, the HMAA four-level authority model with mandatory hysteresis and the FLAME deliberation window directly implement the Directive's principle that autonomous systems must provide operators with the ability to exercise appropriate levels of control. This regulatory bidirectionality defense governance principles informing civilian safety architecture is a deliberate design feature demonstrating pipeline portability across domains.

12. Data Availability

All data, simulation code, engineering artifacts (BOM, electrical/mechanical connection specifications, assembly guide, schematic), and the interactive governance simulator are openly available at DOI: 10.5281/zenodo.19232130 under Creative Commons Attribution 4.0 International (CC BY 4.0). No access restrictions apply. The deposit includes all materials necessary to reproduce the simulation results reported in this paper and to independently evaluate the hardware architecture.

13. Dual-Use Ethics Statement

The BLADE-AV governance pipeline is architecturally identical to the BLADE-EDGE defense variant (DOI: 10.5281/zenodo.19177472). This cross-domain portability is intentional and demonstrates that authority-governed autonomy continuous sensor trust fusion, graded authority with hysteresis, deliberation windows, and deterministic recovery is a domain-agnostic safety principle applicable wherever autonomous systems exercise physical authority. The transfer from defense to civilian transportation does not introduce weapons-enabling capability; the BLADE-AV effector is a normally-open safety relay governing drive-by-wire authority, not a weapons release mechanism. The civilian variant operates under NHTSA ADS voluntary guidance, ISO 26262 functional safety standards, and SAE J3016 automation taxonomy regulatory frameworks that impose independent civilian safety oversight. The author affirms that no export-controlled (ITAR/EAR) technical data is disclosed in this deposit, and that all engineering artifacts are published at a conceptual/architectural level consistent with open academic research.

14. How to Cite

APA

Oktenli, B. (2026). BLADE-AV Governance Node: Authority-Governed Drive-by-Wire Safety Architecture for Autonomous Vehicles (v1.0) [Technical Report]. Georgetown University, MPS Applied Intelligence. https://doi.org/10.5281/zenodo.19232130

BibTeX

```
@techreport{oktenli2026bladeav,   author={Oktenli, Burak},   title={BLADE-AV Governance Node: Authority-Governed Drive-by-Wire Safety Architecture for Autonomous Vehicles},   year={2026},
```

```
version={v1.0},   institution={Georgetown University, MPS Applied Intelligence},
publisher={Zenodo},   doi={10.5281/zenodo.19232130},
url={https://doi.org/10.5281/zenodo.19232130},   license={CC-BY-4.0}}
```

15. References

Note: All references have been individually verified as real, published works.

[1] Oktenli, B. (2026). SATA: A Hardware-Anchored τ-Chain Protocol. Zenodo. https://doi.org/10.5281/zenodo.18936251

[2] Oktenli, B. (2026). HMAA: Hierarchical Mission Authority Architecture. Zenodo. https://doi.org/10.5281/zenodo.18861653

[3] Oktenli, B. (2026). FLAME: Flash War Latency Architecture. Zenodo. https://doi.org/10.5281/zenodo.19015618

[4] Oktenli, B. (2026). CARA: Control Authority Regulation Architecture. Zenodo. https://doi.org/10.5281/zenodo.18917790

[5] Oktenli, B. (2026). BLADE-EDGE Governance Dev Kit (v5.0.3). Zenodo. https://doi.org/10.5281/zenodo.19177472

[6] Shafer, G. (1976). A Mathematical Theory of Evidence. Princeton University Press.

[7] Khaleghi, B. et al. (2013). Multisensor data fusion: A review. Information Fusion, 14(1), 28–44.

[8] U.S. DoD. (2023). Directive 3000.09: Autonomy in Weapon Systems.

[9] NHTSA. (2017). Automated Driving Systems 2.0: A Vision for Safety.

[10] SAE International. (2021). J3016: Driving Automation Systems Taxonomy.

[11] ISO. (2018). ISO 26262: Road vehicles — Functional safety.

[12] Sha, L. et al. (2001). The Simplex Architecture. IEEE Real-Time Systems Symp., 2–10.

[13] Kephart, J. & Chess, D. (2003). Autonomic Computing. IEEE Computer, 36(1), 41–50.

[14] Parasuraman, R. et al. (2000). Human Interaction with Automation. IEEE Trans. SMC-A, 30(3), 286–297.

[15] DARPA. (2024). Assured Autonomy Program. https://www.darpa.mil/program/assured-autonomy

[16] Faul, F. et al. (2007). G*Power 3. Behavior Research Methods, 39(2), 175–191.

[17] Oktenli, B. (2026). ADARA: Adversarial Deception-Aware Risk Architecture. Zenodo. https://doi.org/10.5281/zenodo.19043924

[18] Oktenli, B. (2026). MAIVA: Multi-Agent Integrity Verification Architecture. Zenodo. https://doi.org/10.5281/zenodo.19015517

[19] Shalev-Shwartz, S., Shammah, S., & Shashua, A. (2017). On a Formal Model of Safe and Scalable Self-driving Cars. arXiv:1708.06374.

[20] Nister, D. et al. (2019). Safety Force Field. NVIDIA Technical Report.

[21] ISO. (2022). ISO/PAS 21448: Road vehicles — Safety of the intended functionality (SOTIF).