

SEC-051 - GAISSF Automotive Sector Implementation Guide

Automotive Sector Implementation Guide

Document ID	SEC-051
Status	Publication Candidate - Open Review Gates
Version	1.0
Edition date	1 July 2026

© 2026 ODA3 Pvt Ltd. Published by ODA3 Institute.

Status: Publication Candidate - Open Review Gates

Publisher: ODA3 Institute

Legal entity: ODA3 Pvt Ltd

Edition date: 1 July 2026

Methodology and evidence note. Authoritative GAISSF identifiers and titles are sourced from GAISSF-NOR-001 [T1]. Automotive interpretations, classifications, use cases, threat scenarios, metrics, and crosswalks are analytical implementation guidance [T3] unless a record cites a different tier. No proprietary vehicle telemetry, client data, or internal incident dataset was used.

Executive overview

Automotive artificial intelligence can affect release authority, product liability exposure, warranty and recall cost, type-approval schedules, supplier accountability, vehicle availability, and customer trust. Leadership therefore needs one governed assurance model linking AI security to vehicle cybersecurity, functional safety, Safety of the Intended Functionality (SOTIF), privacy, product safety, software updates, and evidence-based release decisions. [T3]

This guide translates the authoritative 59-control GAISSF baseline into operational automotive contexts. It does not alter the normative controls or claim legal compliance, homologation, safety approval, cybersecurity approval, accredited certification, or prevention of harm. [T1]

Purpose, scope and intended use

The guide applies to vehicle manufacturers, suppliers, semiconductor and software providers, cloud and connected-service providers, fleets, mobility platforms, charging operators, dealerships, aftermarket providers, laboratories, approval functions, and assurance teams. Applicability varies by actor, system boundary, vehicle category, jurisdiction, authority, connectivity, lifecycle stage, and operating design domain. [T3]

Normative and informative boundary

The exact GAISSE control identifiers and titles are authoritative [T1]. Every automotive interpretation, implementation example, evidence record, metric, classification profile, use case, threat scenario, and external cross-standard mapping is informative [T3]. “Shall” is reserved for reproduced normative language.

Automotive ecosystem and system landscape

The relevant system boundary can include perception, sensor fusion, driver and occupant monitoring, vehicle-control functions, battery and energy-management systems, telematics, gateways, infotainment, connected-vehicle backends, mobile applications, public-key infrastructure, mapping, charging, fleet platforms, model-development pipelines, simulation, hardware-in-the-loop environments, manufacturing inspection, and engineering assistants. [T3]

Automotive risk context

Automotive AI risk is differentiated by physical consequence, fleet-scale common-mode failure, long support periods, complex supplier chains, constrained field remediation, signed software and model updates, connected-service dependencies, and the need to reconcile security changes with safety assumptions. [T3]

Risk should be reviewed across six separate impact vectors: Mission, Human, Security, Legal and Policy, Strategic, and Financial. Do not collapse them into one score where aggregation could conceal an unacceptable human or safety impact. A practical review can rate each vector Low, Medium, High, or Critical, record the evidence and uncertainty, and require cross-functional approval where any vector exceeds the organisation’s threshold.

System classification profiles

Profiles A-E are informative implementation aids, not GAISSE conformance levels, vehicle automation levels, Automotive Safety Integrity Levels, assurance levels, maturity levels, or legal classifications. Profile A covers low-impact support; B operational assistance; C material decision support; D safety- or security-relevant automation; and E high-consequence vehicle authority. Classification depends on actual authority, autonomy, safety relevance, connectivity, scale, data sensitivity, reversibility, fallback, supplier dependency, and failure propagation. [T3]

Integrated assurance operating model

1. Define the system, operating context, interfaces, and decision authority.
2. Link AI, cybersecurity, functional-safety, SOTIF, privacy, product-safety, legal, supplier, and product-governance owners.
3. Maintain one controlled inventory and change record.
4. Reconcile assumptions before release and after material change.
5. Preserve evidence sufficient to reconstruct model, data, software, update, vehicle, and fleet state.
6. Escalate safety-relevant security events through both incident-response and product-safety channels.
7. Require explicit residual-risk acceptance and expiry for exceptions. [T3]

90-day implementation sequence

Days 0-30: establish governance; inventory systems; classify authority and consequence; identify critical suppliers; open evidence repositories and specialist review gates.

Days 31-60: complete priority mappings; perform integrated threat analysis; verify model and data provenance; assess suppliers; define monitoring, incident, and Over-the-Air (OTA) update controls.

Days 61-90: test operating effectiveness; execute representative scenarios; validate monitoring and escalation; close critical evidence gaps; conduct independent readiness review; document residual risk.

Worked metric example

Metric: Percentage of applicable in-scope systems with current approved evidence and no overdue high-risk exception.

Example: Five systems are in scope. Four have current approved evidence and no overdue high-risk exception. The result is $4 \div 5 = 80\%$. The report must retain the denominator, exclusions, evidence date, and exception status so the figure cannot be improved by silently narrowing scope.

Standards and regulatory considerations

Mappings are directional and do not establish equivalence or compliance inheritance. Primary official publications are treated as [T1]. The automotive interpretation and crosswalk are analytical [T3] until clause-level specialist review is completed.

- **UN Regulation No. 155 - Cyber security and cyber security management system** [T1]: Vehicle cybersecurity governance, lifecycle risk management, monitoring and response. Analytical mapping [T3]. Limitation: Jurisdiction and vehicle-category applicability require legal/type-approval review.
- **UN Regulation No. 156 - Software update and software update management system** [T1]: Software/AI update governance, integrity, records, compatibility and controlled release. Analytical mapping [T3]. Limitation: Current series of amendments and national implementation must be verified.
- **ISO/SAE 21434:2021 - Road vehicles - Cybersecurity engineering** [T1]: Cybersecurity engineering across concept, development, production, operations, maintenance and decommissioning. Analytical mapping [T3]. Limitation: No equivalence or compliance inheritance is claimed.
- **ISO 26262 series:2018 - Road vehicles - Functional safety** [T1]: Functional-safety management and lifecycle coordination for safety-related E/E systems. Analytical mapping [T3]. Limitation: AI security does not replace functional-safety analysis or confirmation measures.
- **ISO 21448:2022 - Road vehicles - Safety of the intended functionality** [T1]: Performance limitations, foreseeable misuse, scenario analysis and unknown unsafe conditions. Analytical mapping [T3]. Limitation: Clause-level mapping requires lawful access and specialist SOTIF review.
- **ISO 24089:2023 - Road vehicles - Software update engineering** [T1]: Software update engineering and coordination with AI/model releases. Analytical mapping [T3]. Limitation: Applicability to model-only changes must be determined in system context.
- **ISO/IEC 42001:2023 - Information technology - Artificial intelligence - Management system** [T1]: AI governance, objectives, controls, documented information and continual improvement. Analytical mapping [T3]. Limitation: Management-system conformity does not establish vehicle safety or cybersecurity.

- **NIST AI RMF 1.0 - Artificial Intelligence Risk Management Framework [T1]:** Govern, map, measure and manage AI risks. Analytical mapping [T3]. Limitation: Not automotive-specific and not a legal compliance instrument.

Domain-by-domain automotive implementation guidance

D1 - MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS

Apply this domain through data provenance, model integrity, adversarial testing, drift monitoring, signing, conversion and release controls. [T3] The Control Master and Annex A define the exact automotive implementation records.

D2 - RUNTIME SECURITY & ADVERSARIAL DEFENSE

Apply this domain through prompt, multimodal, tool-call and cross-context controls for generative and interactive automotive functions. [T3] The Control Master and Annex A define the exact automotive implementation records.

D3 - AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY

Apply this domain through authority limits, agent communication, memory, embodied action, fallback, oversight and command verification. [T3] The Control Master and Annex A define the exact automotive implementation records.

D4 - SUPPLY CHAIN & THIRD-PARTY AI SECURITY

Apply this domain through AI bills of materials, model scanning, registry vetting, third-party APIs, shadow AI and software composition analysis. [T3] The Control Master and Annex A define the exact automotive implementation records.

D5 - CONTENT SAFETY & OUTPUT INTEGRITY

Apply this domain through output integrity, privacy, personal information leakage, copyright and privacy-preserving machine learning. [T3] The Control Master and Annex A define the exact automotive implementation records.

D6 - GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT

Apply this domain through human oversight, auditability, model cards, incident response, decommissioning, supplier governance and resilience. [T3] The Control Master and Annex A define the exact automotive implementation records.

D7 - HUMAN & SOCIETAL HARMS

Apply this domain through training, authentication, social-engineering response and AI-enhanced external-attack defence. [T3] The Control Master and Annex A define the exact automotive implementation records.

D8 - REGULATORY ALIGNMENT & COMPLIANCE

Apply this domain through applicable-law mapping, management-system alignment, technical documentation and secure-development evidence. [T3] The Control Master and Annex A define the exact automotive implementation records.

D9 - PHYSICAL AI SAFETY

Apply this domain through physical boundaries, safe states, override, cyber-physical detection, environmental monitoring and evidence preservation. [T3] The Control Master and Annex A define the exact automotive implementation records.

Annex A - Complete control mapping

D1-CTL-01 - DATASET PROVENANCE & POISONING PREVENTION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Dataset Provenance & Poisoning Prevention to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: ADAS/automated-driving function; Connected-vehicle backend

Related use cases: AUC-028; AUC-029

Related threat scenarios: None assigned

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-005; AEV-009; AEV-015

Assessment question: Is dataset provenance & poisoning prevention implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-02 - MODEL EXTRACTION RESISTANCE

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Model Extraction Resistance to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Driver or occupant monitoring; OTA update infrastructure

Related use cases: AUC-028; AUC-029; AUC-030

Related threat scenarios: ATS-001

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-015

Assessment question: Is model extraction resistance implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-03 - BEHAVIORAL DRIFT DETECTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Behavioral Drift Detection to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: In-vehicle AI component; Fleet analytics platform

Related use cases: AUC-001; AUC-029; AUC-030

Related threat scenarios: ATS-002

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-005; AEV-007; AEV-010; AEV-014; AEV-016; AEV-019; AEV-020; AEV-025

Assessment question: Is behavioral drift detection implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-04 - FEDERATED LEARNING POISONING PREVENTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Federated Learning Poisoning Prevention to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Connected-vehicle backend; Engineering AI assistant

Related use cases: AUC-001; AUC-029; AUC-030; AUC-031

Related threat scenarios: ATS-003

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-009

Assessment question: Is federated learning poisoning prevention implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-05 - EMBEDDING SPACE ROBUSTNESS

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Embedding Space Robustness to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: OTA update infrastructure; Manufacturing AI system

Related use cases: AUC-001; AUC-002; AUC-030; AUC-031

Related threat scenarios: ATS-004

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-014

Assessment question: Is embedding space robustness implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-06 - POST-QUANTUM MODEL SIGNING & CRYPTO HARDENING

Authoritative GAISF requirement [T1]: Refer to GAISF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Post-Quantum Model Signing & Crypto Hardening to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Fleet analytics platform; Battery and energy-management AI

Related use cases: AUC-001; AUC-002; AUC-030; AUC-031; AUC-032

Related threat scenarios: ATS-005

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-017; AEV-018

Assessment question: Is post-quantum model signing & crypto hardening implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-07 - LORA/ADAPTER INTEGRITY VERIFICATION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Lora/Adapter Integrity Verification to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Engineering AI assistant; Vehicle security monitoring

Related use cases: AUC-001; AUC-002; AUC-003; AUC-031; AUC-032

Related threat scenarios: ATS-006

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-018

Assessment question: Is lora/adapter integrity verification implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-08 - MODEL MERGE ATTACK DETECTION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Model Merge Attack Detection to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Manufacturing AI system; ADAS/automated-driving function

Related use cases: AUC-002; AUC-003; AUC-031; AUC-032

Related threat scenarios: ATS-007

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-018

Assessment question: Is model merge attack detection implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D1-CTL-09 - QUANTIZATION BACKDOOR SCREENING

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Quantization Backdoor Screening to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Battery and energy-management AI; Driver or occupant monitoring

Related use cases: AUC-002; AUC-003; AUC-004; AUC-032

Related threat scenarios: ATS-001; ATS-008

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-018

Assessment question: Is quantization backdoor screening implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D2-CTL-01 - DIRECT PROMPT INJECTION PREVENTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Direct Prompt Injection Prevention to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Vehicle security monitoring; In-vehicle AI component

Related use cases: AUC-003; AUC-004; AUC-032

Related threat scenarios: ATS-002; ATS-009

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-005; AEV-008; AEV-015

Assessment question: Is direct prompt injection prevention implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D2-CTL-02 - INDIRECT PROMPT INJECTION PREVENTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Indirect Prompt Injection Prevention to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: ADAS/automated-driving function; Connected-vehicle backend

Related use cases: AUC-003; AUC-004; AUC-005

Related threat scenarios: ATS-003; ATS-010

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-005; AEV-008; AEV-015

Assessment question: Is indirect prompt injection prevention implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D2-CTL-03 - JAILBREAK RESISTANCE TESTING

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Jailbreak Resistance Testing to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Driver or occupant monitoring; OTA update infrastructure

Related use cases: AUC-004; AUC-005

Related threat scenarios: ATS-004; ATS-011

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-014; AEV-015

Assessment question: Is jailbreak resistance testing implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D2-CTL-04 - MULTI-MODAL INJECTION DEFENSE

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Multi-Modal Injection Defense to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: In-vehicle AI component; Fleet analytics platform

Related use cases: AUC-004; AUC-005; AUC-006

Related threat scenarios: ATS-005; ATS-012

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-008; AEV-014; AEV-015

Assessment question: Is multi-modal injection defense implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D2-CTL-05 - FUNCTION CALL/TOOL CALL INJECTION PREVENTION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Function Call/Tool Call Injection Prevention to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Connected-vehicle backend; Engineering AI assistant

Related use cases: AUC-005; AUC-006

Related threat scenarios: ATS-006; ATS-013

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-004; AEV-015; AEV-017

Assessment question: Is function call/tool call injection prevention implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D2-CTL-06 - CROSS-CONTEXT HIJACKING MITIGATION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Cross-Context Hijacking Mitigation to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: OTA update infrastructure; Manufacturing AI system

Related use cases: AUC-005; AUC-006; AUC-007

Related threat scenarios: ATS-007; ATS-014

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-004; AEV-015

Assessment question: Is cross-context hijacking mitigation implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D3-CTL-01 - LEAST AGENCY ENFORCEMENT

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Least Agency Enforcement to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Fleet analytics platform; Battery and energy-management AI

Related use cases: AUC-006; AUC-007

Related threat scenarios: ATS-001; ATS-008; ATS-015

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-002; AEV-004; AEV-017

Assessment question: Is least agency enforcement implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D3-CTL-02 - INTER-AGENT COMMUNICATION SECURITY

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Inter-Agent Communication Security to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Engineering AI assistant; Vehicle security monitoring

Related use cases: AUC-006; AUC-007; AUC-008

Related threat scenarios: ATS-002; ATS-009; ATS-016

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-002; AEV-004; AEV-008

Assessment question: Is inter-agent communication security implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D3-CTL-03 - AGENTIC PROMPT CHAINING DETECTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Agentic Prompt Chaining Detection to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Manufacturing AI system; ADAS/automated-driving function

Related use cases: AUC-007; AUC-008

Related threat scenarios: ATS-003; ATS-010; ATS-017

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is agentic prompt chaining detection implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D3-CTL-04 - EMBODIED AI SAFETY CONTROLS

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Embodied Ai Safety Controls to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Battery and energy-management AI; Driver or occupant monitoring

Related use cases: AUC-007; AUC-008; AUC-009

Related threat scenarios: ATS-004; ATS-011; ATS-018

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-002; AEV-005; AEV-006; AEV-007; AEV-016

Assessment question: Is embodied ai safety controls implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D3-CTL-05 - MULTI-AGENT TRUST CHAIN ATTESTATION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Multi-Agent Trust Chain Attestation to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Vehicle security monitoring; In-vehicle AI component

Related use cases: AUC-008; AUC-009

Related threat scenarios: ATS-005; ATS-012; ATS-019

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is multi-agent trust chain attestation implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D3-CTL-06 - PERSISTENT MEMORY EXFILTRATION PREVENTION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Persistent Memory Exfiltration Prevention to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: ADAS/automated-driving function; Connected-vehicle backend

Related use cases: AUC-008; AUC-009; AUC-010

Related threat scenarios: ATS-006; ATS-013; ATS-020

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is persistent memory exfiltration prevention implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D3-CTL-07 - SECURE MEMORY LIFECYCLE MANAGEMENT

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Secure Memory Lifecycle Management to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Driver or occupant monitoring; OTA update infrastructure

Related use cases: AUC-009; AUC-010

Related threat scenarios: ATS-007; ATS-014; ATS-021

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is secure memory lifecycle management implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D4-CTL-01 - AI BILL OF MATERIALS (AI BOM) MAINTENANCE

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai Bill Of Materials (Ai Bom) Maintenance to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: In-vehicle AI component; Fleet analytics platform

Related use cases: AUC-009; AUC-010; AUC-011

Related threat scenarios: ATS-001; ATS-008; ATS-015; ATS-022

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-012; AEV-018

Assessment question: Is ai bill of materials (ai bom) maintenance implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D4-CTL-02 - MODEL FILE & ARTIFACT SCANNING

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Model File & Artifact Scanning to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Connected-vehicle backend; Engineering AI assistant

Related use cases: AUC-010; AUC-011

Related threat scenarios: ATS-002; ATS-009; ATS-016; ATS-023

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-012; AEV-017; AEV-018; AEV-025

Assessment question: Is model file & artifact scanning implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D4-CTL-03 - MODEL HUB & REGISTRY VETTING

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Model Hub & Registry Vetting to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: OTA update infrastructure; Manufacturing AI system

Related use cases: AUC-010; AUC-011; AUC-012

Related threat scenarios: ATS-003; ATS-010; ATS-017; ATS-024

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-012

Assessment question: Is model hub & registry vetting implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D4-CTL-04 - MCP SERVER BEHAVIORAL MONITORING

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Mcp Server Behavioral Monitoring to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Fleet analytics platform; Battery and energy-management AI

Related use cases: AUC-011; AUC-012

Related threat scenarios: ATS-004; ATS-011; ATS-018; ATS-025

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-019; AEV-020

Assessment question: Is mcp server behavioral monitoring implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D4-CTL-05 - THIRD-PARTY AI API SECURITY ASSESSMENT

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Third-Party Ai Api Security Assessment to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Engineering AI assistant; Vehicle security monitoring

Related use cases: AUC-011; AUC-012; AUC-013

Related threat scenarios: ATS-005; ATS-012; ATS-019; ATS-026

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-004; AEV-005; AEV-008; AEV-012; AEV-013

Assessment question: Is third-party ai api security assessment implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D4-CTL-06 - SHADOW AI DISCOVERY & GOVERNANCE

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Shadow Ai Discovery & Governance to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Manufacturing AI system; ADAS/automated-driving function

Related use cases: AUC-012; AUC-013

Related threat scenarios: ATS-006; ATS-013; ATS-020; ATS-027

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is shadow ai discovery & governance implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D4-CTL-07 - AI SOFTWARE COMPOSITION ANALYSIS (SCA)

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai Software Composition Analysis (Sca) to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Battery and energy-management AI; Driver or occupant monitoring

Related use cases: AUC-012; AUC-013; AUC-014

Related threat scenarios: ATS-007; ATS-014; ATS-021; ATS-028

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-011; AEV-012; AEV-017

Assessment question: Is ai software composition analysis (sca) implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D5-CTL-01 - HARMFUL CONTENT BLOCKING

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Harmful Content Blocking to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Vehicle security monitoring; In-vehicle AI component

Related use cases: AUC-013; AUC-014

Related threat scenarios: ATS-008; ATS-015; ATS-022; ATS-029

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is harmful content blocking implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D5-CTL-02 - PII LEAKAGE PREVENTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Pii Leakage Prevention to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: ADAS/automated-driving function; Connected-vehicle backend

Related use cases: AUC-013; AUC-014; AUC-015

Related threat scenarios: ATS-009; ATS-016; ATS-023; ATS-030

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-009

Assessment question: Is pii leakage prevention implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D5-CTL-03 - COPYRIGHT DETECTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Copyright Detection to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Driver or occupant monitoring; OTA update infrastructure

Related use cases: AUC-014; AUC-015

Related threat scenarios: ATS-010; ATS-017; ATS-024; ATS-031

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is copyright detection implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D5-CTL-04 - AI WATERMARKING ROBUSTNESS

Authoritative GAISF requirement [T1]: Refer to GAISF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai Watermarking Robustness to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: In-vehicle AI component; Fleet analytics platform

Related use cases: AUC-014; AUC-015; AUC-016

Related threat scenarios: ATS-011; ATS-018; ATS-025; ATS-032

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is ai watermarking robustness implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D5-CTL-05 - PRIVACY-BY-DESIGN VERIFICATION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Privacy-By-Design Verification to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Connected-vehicle backend; Engineering AI assistant

Related use cases: AUC-015; AUC-016

Related threat scenarios: ATS-012; ATS-019; ATS-026; ATS-033

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-009

Assessment question: Is privacy-by-design verification implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D5-CTL-06 - PRIVACY-PRESERVING ML VALIDATION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Privacy-Preserving ML Validation to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: OTA update infrastructure; Manufacturing AI system

Related use cases: AUC-015; AUC-016; AUC-017

Related threat scenarios: ATS-013; ATS-020; ATS-027; ATS-034

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-009

Assessment question: Is privacy-preserving ml validation implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D6-CTL-01 - HUMAN-IN-THE-LOOP FOR HIGH-RISK ACTIONS

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Human-In-The-Loop For High-Risk Actions to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Fleet analytics platform; Battery and energy-management AI

Related use cases: AUC-016; AUC-017

Related threat scenarios: ATS-014; ATS-021; ATS-028

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-003; AEV-006; AEV-007; AEV-022; AEV-023

Assessment question: Is human-in-the-loop for high-risk actions implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D6-CTL-02 - AUDIT TRAIL COMPLETENESS

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Audit Trail Completeness to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Engineering AI assistant; Vehicle security monitoring

Related use cases: AUC-016; AUC-017; AUC-018

Related threat scenarios: ATS-015; ATS-022; ATS-029

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-001; AEV-003; AEV-019; AEV-021; AEV-022; AEV-024; AEV-025

Assessment question: Is audit trail completeness implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D6-CTL-03 - AI MODEL CARD COMPLETENESS

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai Model Card Completeness to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Manufacturing AI system; ADAS/automated-driving function

Related use cases: AUC-017; AUC-018

Related threat scenarios: ATS-016; ATS-023; ATS-030

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-010

Assessment question: Is ai model card completeness implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D6-CTL-04 - AI INCIDENT RESPONSE READINESS

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai Incident Response Readiness to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Battery and energy-management AI; Driver or occupant monitoring

Related use cases: AUC-017; AUC-018; AUC-019

Related threat scenarios: ATS-017; ATS-024; ATS-031

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-005; AEV-019; AEV-020; AEV-021; AEV-023; AEV-024; AEV-025

Assessment question: Is ai incident response readiness implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D6-CTL-05 - MODEL DEPRECIATION & DECOMMISSIONING

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Model Deprecation & Decommissioning to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Vehicle security monitoring; In-vehicle AI component

Related use cases: AUC-018; AUC-019

Related threat scenarios: ATS-018; ATS-025; ATS-032

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is model depreciation & decommissioning implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D6-CTL-06 - THIRD-PARTY AI VENDOR GOVERNANCE

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Third-Party Ai Vendor Governance to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: ADAS/automated-driving function; Connected-vehicle backend

Related use cases: AUC-018; AUC-019; AUC-020

Related threat scenarios: ATS-019; ATS-026; ATS-033

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-001; AEV-012; AEV-013; AEV-022; AEV-023; AEV-024

Assessment question: Is third-party ai vendor governance implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D6-CTL-07 - AI RESILIENCE & BUSINESS CONTINUITY

Authoritative GAISF requirement [T1]: Refer to GAISF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai Resilience & Business Continuity to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Driver or occupant monitoring; OTA update infrastructure

Related use cases: AUC-019; AUC-020

Related threat scenarios: ATS-020; ATS-027; ATS-034

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-019; AEV-020; AEV-021; AEV-022; AEV-024; AEV-025

Assessment question: Is ai resilience & business continuity implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D7-CTL-H01 - AI-GENERATED PHISHING SIMULATION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai-Generated Phishing Simulation to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: In-vehicle AI component; Fleet analytics platform

Related use cases: AUC-019; AUC-020; AUC-021

Related threat scenarios: ATS-021; ATS-028

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-023

Assessment question: Is ai-generated phishing simulation implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D7-CTL-H02 - DEEPPFAKE DETECTION TRAINING

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Deepfake Detection Training to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Connected-vehicle backend; Engineering AI assistant

Related use cases: AUC-020; AUC-021

Related threat scenarios: ATS-022; ATS-029

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-023

Assessment question: Is deepfake detection training implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D7-CTL-H03 - OUT-OF-BAND AUTHENTICATION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Out-Of-Band Authentication to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: OTA update infrastructure; Manufacturing AI system

Related use cases: AUC-020; AUC-021; AUC-022

Related threat scenarios: ATS-023; ATS-030

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-023

Assessment question: Is out-of-band authentication implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D7-CTL-H04 - AI SOCIAL ENGINEERING IR

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai Social Engineering Ir to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Fleet analytics platform; Battery and energy-management AI

Related use cases: AUC-021; AUC-022

Related threat scenarios: ATS-024; ATS-031

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-021; AEV-023

Assessment question: Is ai social engineering ir implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D7-CTL-H05 - AI-ENHANCED EXTERNAL ATTACK DEFENSE

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Ai-Enhanced External Attack Defense to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Engineering AI assistant; Vehicle security monitoring

Related use cases: AUC-021; AUC-022; AUC-023

Related threat scenarios: ATS-025; ATS-032

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is ai-enhanced external attack defense implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D8-CTL-01 - EU AI ACT RISK TIER MAPPING

Authoritative GAISF requirement [T1]: Refer to GAISF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Eu Ai Act Risk Tier Mapping to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Manufacturing AI system; ADAS/automated-driving function

Related use cases: AUC-022; AUC-023

Related threat scenarios: ATS-026; ATS-033

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-001; AEV-003; AEV-013; AEV-022; AEV-024

Assessment question: Is eu ai act risk tier mapping implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D8-CTL-02 - ISO 42001 GAP ANALYSIS

Authoritative GAISF requirement [T1]: Refer to GAISF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Iso 42001 Gap Analysis to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Battery and energy-management AI; Driver or occupant monitoring

Related use cases: AUC-022; AUC-023; AUC-024

Related threat scenarios: ATS-027; ATS-034

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-001; AEV-003; AEV-024; AEV-025

Assessment question: Is iso 42001 gap analysis implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D8-CTL-03 - GPAI TECHNICAL DOCUMENTATION VERIFICATION

Authoritative GAISF requirement [T1]: Refer to GAISF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Gpai Technical Documentation Verification to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Vehicle security monitoring; In-vehicle AI component

Related use cases: AUC-023; AUC-024

Related threat scenarios: ATS-028

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-010; AEV-013

Assessment question: Is gpai technical documentation verification implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D8-CTL-04 - DORA ICT INCIDENT REPORTING (FINANCIAL SECTOR)

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Dora Ict Incident Reporting (Financial Sector) to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: ADAS/automated-driving function; Connected-vehicle backend

Related use cases: AUC-023; AUC-024; AUC-025

Related threat scenarios: ATS-029

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: To be determined

Assessment question: Is dora ict incident reporting (financial sector) implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D8-CTL-05 - NIST SP 800-218A COMPLIANCE CHECK

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Nist Sp 800-218A Compliance Check to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Driver or occupant monitoring; OTA update infrastructure

Related use cases: AUC-024; AUC-025

Related threat scenarios: ATS-030

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-013; AEV-017; AEV-025

Assessment question: Is nist sp 800-218a compliance check implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D9-CTL-01 - PHYSICAL HARM BOUNDARY ENFORCEMENT

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Physical Harm Boundary Enforcement to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: In-vehicle AI component; Fleet analytics platform

Related use cases: AUC-024; AUC-025; AUC-026

Related threat scenarios: ATS-031

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-002; AEV-003; AEV-006; AEV-007; AEV-016; AEV-022

Assessment question: Is physical harm boundary enforcement implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D9-CTL-02 - SAFE STATE AND GRACEFUL DEGRADATION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Safe State And Graceful Degradation to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Connected-vehicle backend; Engineering AI assistant

Related use cases: AUC-025; AUC-026

Related threat scenarios: ATS-032

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-002; AEV-006; AEV-007; AEV-014; AEV-016

Assessment question: Is safe state and graceful degradation implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Data Governance Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D9-CTL-03 - HUMAN OVERRIDE AND EMERGENCY STOP

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Human Override And Emergency Stop to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: OTA update infrastructure; Manufacturing AI system

Related use cases: AUC-025; AUC-026; AUC-027

Related threat scenarios: ATS-033

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-006

Assessment question: Is human override and emergency stop implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: AI Security Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D9-CTL-04 - CYBER-PHYSICAL ATTACK DETECTION

Authoritative GAISSE requirement [T1]: Refer to GAISSE-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Cyber-Physical Attack Detection to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Fleet analytics platform; Battery and energy-management AI

Related use cases: AUC-026; AUC-027

Related threat scenarios: ATS-034

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-005; AEV-008; AEV-015; AEV-019; AEV-020; AEV-021

Assessment question: Is cyber-physical attack detection implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Automotive Cybersecurity Lead / Privacy Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D9-CTL-05 - PHYSICAL ENVIRONMENT INTEGRITY MONITORING

Authoritative GAISF requirement [T1]: Refer to GAISF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Physical Environment Integrity Monitoring to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Engineering AI assistant; Vehicle security monitoring

Related use cases: AUC-026; AUC-027; AUC-028

Related threat scenarios: None assigned

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-007; AEV-014; AEV-016; AEV-019; AEV-020

Assessment question: Is physical environment integrity monitoring implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Vehicle Software Lead / Independent Assessor

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D9-CTL-06 - ACTUATOR COMMAND VERIFICATION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Actuator Command Verification to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Manufacturing AI system; ADAS/automated-driving function

Related use cases: AUC-027; AUC-028

Related threat scenarios: None assigned

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-002; AEV-004; AEV-006; AEV-008; AEV-017; AEV-018

Assessment question: Is actuator command verification implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Functional Safety Lead / Product Security Lead

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

D9-CTL-07 - PHYSICAL INCIDENT EVIDENCE PRESERVATION

Authoritative GAISSF requirement [T1]: Refer to GAISSF-NOR-001 v1.0.

Automotive interpretation [T3]: Apply Physical Incident Evidence Preservation to automotive AI components and their vehicle, backend, engineering, manufacturing, supplier, update, and monitoring dependencies using a documented risk-based scope.

Applicability: Determine for every in-scope system; record Not Applicable only with approved rationale.

Representative systems: Battery and energy-management AI; Driver or occupant monitoring

Related use cases: AUC-027; AUC-028; AUC-029

Related threat scenarios: None assigned

Minimum credible implementation: Approved scope and owner; documented applicability; baseline technical or procedural safeguard; retained evidence; defined exception and escalation route.

Enhanced implementation: Automated enforcement and telemetry; independent testing; fleet correlation where relevant; supplier attestations backed by evidence; change-triggered reassessment.

Illustrative evidence: AEV-006; AEV-020; AEV-021; AEV-025

Assessment question: Is physical incident evidence preservation implemented across the approved automotive system boundary, including suppliers, updates, field monitoring, exceptions, and operating evidence?

Owner / reviewer: Supplier Assurance Lead / Functional Safety Reviewer

Safety coordination: Reconcile security assumptions and mitigations with functional-safety and SOTIF artefacts; do not alter safety mechanisms without impact analysis.

Notably absent: No claim of automatic legal compliance, homologation, functional-safety approval, cybersecurity approval, or prevention of harm.

Review status: Open - automotive specialist review required

Annex B - Automotive AI use-case catalogue

AUC-001 - Object and pedestrian detection

Context: Automotive AI application supporting object and pedestrian detection within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-1 supplier

Authority / safety relevance: High / High

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D1-CTL-03; D1-CTL-04; D1-CTL-05; D1-CTL-06; D1-CTL-07

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-002 - Lane detection and lane keeping

Context: Automotive AI application supporting lane detection and lane keeping within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-2 supplier

Authority / safety relevance: High / High

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D1-CTL-05; D1-CTL-06; D1-CTL-07; D1-CTL-08; D1-CTL-09

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-003 - Traffic-sign recognition

Context: Automotive AI application supporting traffic-sign recognition within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Automotive semiconductor supplier

Authority / safety relevance: Medium / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D1-CTL-07; D1-CTL-08; D1-CTL-09; D2-CTL-01; D2-CTL-02

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-004 - Driver-state monitoring

Context: Automotive AI application supporting driver-state monitoring within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Vehicle software platform provider

Authority / safety relevance: Low / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D1-CTL-09; D2-CTL-01; D2-CTL-02; D2-CTL-03; D2-CTL-04

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-005 - Occupant classification

Context: Automotive AI application supporting occupant classification within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Cloud or backend provider

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D2-CTL-02; D2-CTL-03; D2-CTL-04; D2-CTL-05; D2-CTL-06

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-006 - Automated emergency braking

Context: Automotive AI application supporting automated emergency braking within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Fleet operator

Authority / safety relevance: High / High

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D2-CTL-04; D2-CTL-05; D2-CTL-06; D3-CTL-01; D3-CTL-02

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-007 - Path planning and motion prediction

Context: Automotive AI application supporting path planning and motion prediction within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Mobility platform

Authority / safety relevance: High / High

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D2-CTL-06; D3-CTL-01; D3-CTL-02; D3-CTL-03; D3-CTL-04

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-008 - Parking assistance

Context: Automotive AI application supporting parking assistance within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Charging infrastructure operator

Authority / safety relevance: Low / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D3-CTL-02; D3-CTL-03; D3-CTL-04; D3-CTL-05; D3-CTL-06

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-009 - Route optimisation

Context: Automotive AI application supporting route optimisation within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Aftermarket provider

Authority / safety relevance: Medium / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D3-CTL-04; D3-CTL-05; D3-CTL-06; D3-CTL-07; D4-CTL-01

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-010 - Battery state-of-health prediction

Context: Automotive AI application supporting battery state-of-health prediction within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Vehicle OEM

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D3-CTL-06; D3-CTL-07; D4-CTL-01; D4-CTL-02; D4-CTL-03

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-011 - Battery thermal-risk detection

Context: Automotive AI application supporting battery thermal-risk detection within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-1 supplier

Authority / safety relevance: High / High

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D4-CTL-01; D4-CTL-02; D4-CTL-03; D4-CTL-04; D4-CTL-05

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-012 - Predictive maintenance

Context: Automotive AI application supporting predictive maintenance within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-2 supplier

Authority / safety relevance: Medium / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D4-CTL-03; D4-CTL-04; D4-CTL-05; D4-CTL-06; D4-CTL-07

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-013 - Vehicle-network anomaly detection

Context: Automotive AI application supporting vehicle-network anomaly detection within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Automotive semiconductor supplier

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D4-CTL-05; D4-CTL-06; D4-CTL-07; D5-CTL-01; D5-CTL-02

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-014 - Remote diagnostics

Context: Automotive AI application supporting remote diagnostics within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Vehicle software platform provider

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D4-CTL-07; D5-CTL-01; D5-CTL-02; D5-CTL-03; D5-CTL-04

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-015 - Vehicle identity verification

Context: Automotive AI application supporting vehicle identity verification within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Cloud or backend provider

Authority / safety relevance: Medium / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D5-CTL-02; D5-CTL-03; D5-CTL-04; D5-CTL-05; D5-CTL-06

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-016 - Fleet-driver scoring

Context: Automotive AI application supporting fleet-driver scoring within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Fleet operator

Authority / safety relevance: Low / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D5-CTL-04; D5-CTL-05; D5-CTL-06; D6-CTL-01; D6-CTL-02

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-017 - Insurance telematics

Context: Automotive AI application supporting insurance telematics within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Mobility platform

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D5-CTL-06; D6-CTL-01; D6-CTL-02; D6-CTL-03; D6-CTL-04

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-018 - Charging optimisation

Context: Automotive AI application supporting charging optimisation within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Charging infrastructure operator

Authority / safety relevance: Medium / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D6-CTL-02; D6-CTL-03; D6-CTL-04; D6-CTL-05; D6-CTL-06

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-019 - Manufacturing quality inspection

Context: Automotive AI application supporting manufacturing quality inspection within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Aftermarket provider

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D6-CTL-04; D6-CTL-05; D6-CTL-06; D6-CTL-07; D7-CTL-H01

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-020 - Robotic process control

Context: Automotive AI application supporting robotic process control within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Vehicle OEM

Authority / safety relevance: Low / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D6-CTL-06; D6-CTL-07; D7-CTL-H01; D7-CTL-H02; D7-CTL-H03

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-021 - Supplier-risk analytics

Context: Automotive AI application supporting supplier-risk analytics within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-1 supplier

Authority / safety relevance: Medium / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D7-CTL-H01; D7-CTL-H02; D7-CTL-H03; D7-CTL-H04; D7-CTL-H05

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-022 - Software-defect prediction

Context: Automotive AI application supporting software-defect prediction within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-2 supplier

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D7-CTL-H03; D7-CTL-H04; D7-CTL-H05; D8-CTL-01; D8-CTL-02

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-023 - Vulnerability prioritisation

Context: Automotive AI application supporting vulnerability prioritisation within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Automotive semiconductor supplier

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D7-CTL-H05; D8-CTL-01; D8-CTL-02; D8-CTL-03; D8-CTL-04

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-024 - Incident triage

Context: Automotive AI application supporting incident triage within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Vehicle software platform provider

Authority / safety relevance: Medium / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D8-CTL-02; D8-CTL-03; D8-CTL-04; D8-CTL-05; D9-CTL-01

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-025 - Simulation and scenario generation

Context: Automotive AI application supporting simulation and scenario generation within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Cloud or backend provider

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D8-CTL-04; D8-CTL-05; D9-CTL-01; D9-CTL-02; D9-CTL-03

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-026 - Synthetic sensor-data generation

Context: Automotive AI application supporting synthetic sensor-data generation within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Fleet operator

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D9-CTL-01; D9-CTL-02; D9-CTL-03; D9-CTL-04; D9-CTL-05

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-027 - Engineering coding assistant

Context: Automotive AI application supporting engineering coding assistant within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Mobility platform

Authority / safety relevance: Medium / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D9-CTL-03; D9-CTL-04; D9-CTL-05; D9-CTL-06; D9-CTL-07

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-028 - Requirements analysis assistant

Context: Automotive AI application supporting requirements analysis assistant within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Charging infrastructure operator

Authority / safety relevance: Low / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D9-CTL-05; D9-CTL-06; D9-CTL-07; D1-CTL-01; D1-CTL-02

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-029 - Test-case generation

Context: Automotive AI application supporting test-case generation within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Aftermarket provider

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D9-CTL-07; D1-CTL-01; D1-CTL-02; D1-CTL-03; D1-CTL-04

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-030 - Safety-case evidence assistant

Context: Automotive AI application supporting safety-case evidence assistant within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Vehicle OEM

Authority / safety relevance: Medium / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D1-CTL-02; D1-CTL-03; D1-CTL-04; D1-CTL-05; D1-CTL-06

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-031 - Compliance evidence analysis

Context: Automotive AI application supporting compliance evidence analysis within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-1 supplier

Authority / safety relevance: Low / Low

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D1-CTL-04; D1-CTL-05; D1-CTL-06; D1-CTL-07; D1-CTL-08

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

AUC-032 - Customer-service assistant

Context: Automotive AI application supporting customer-service assistant within a defined vehicle, engineering, manufacturing, fleet, or service boundary.

Actors: Tier-2 supplier

Authority / safety relevance: Low / Medium

Failure modes: Incorrect inference, distribution shift, unavailable service, unsafe automation, stale data, or inadequate human escalation.

Applicable controls: D1-CTL-06; D1-CTL-07; D1-CTL-08; D1-CTL-09; D2-CTL-01

Evidence expectations: Approved scope, architecture, threat model, data/model provenance, test results, release decision, monitoring records, and exception log.

Limitations: Classification depends on actual authority, operating design domain, connectivity, scale, and fallback design.

Annex C - Threat and failure scenario catalogue

ATS-001 - Adversarial perception manipulation [T3]

Description: A threat or failure scenario involving adversarial perception manipulation in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-02; D1-CTL-09; D3-CTL-01; D4-CTL-01

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-002 - Sensor spoofing or blinding [T3]

Description: A threat or failure scenario involving sensor spoofing or blinding in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-03; D2-CTL-01; D3-CTL-02; D4-CTL-02

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-003 - Training-data poisoning [T3]

Description: A threat or failure scenario involving training-data poisoning in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-04; D2-CTL-02; D3-CTL-03; D4-CTL-03

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-004 - Compromised simulation data [T3]

Description: A threat or failure scenario involving compromised simulation data in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-05; D2-CTL-03; D3-CTL-04; D4-CTL-04

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-005 - Mapping-data compromise [T3]

Description: A threat or failure scenario involving mapping-data compromise in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-06; D2-CTL-04; D3-CTL-05; D4-CTL-05

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-006 - Model tampering or substitution [T3]

Description: A threat or failure scenario involving model tampering or substitution in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-07; D2-CTL-05; D3-CTL-06; D4-CTL-06

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-007 - Insecure model update [T3]

Description: A threat or failure scenario involving insecure model update in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-08; D2-CTL-06; D3-CTL-07; D4-CTL-07

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-008 - Rollback attack [T3]

Description: A threat or failure scenario involving rollback attack in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D1-CTL-09; D3-CTL-01; D4-CTL-01; D5-CTL-01

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-009 - Unauthorised calibration change [T3]

Description: A threat or failure scenario involving unauthorised calibration change in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D2-CTL-01; D3-CTL-02; D4-CTL-02; D5-CTL-02

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-010 - Third-party AI component compromise [T3]

Description: A threat or failure scenario involving third-party ai component compromise in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D2-CTL-02; D3-CTL-03; D4-CTL-03; D5-CTL-03

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-011 - Unsafe model compression [T3]

Description: A threat or failure scenario involving unsafe model compression in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D2-CTL-03; D3-CTL-04; D4-CTL-04; D5-CTL-04

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-012 - Loss of model provenance [T3]

Description: A threat or failure scenario involving loss of model provenance in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D2-CTL-04; D3-CTL-05; D4-CTL-05; D5-CTL-05

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-013 - Prompt injection in engineering copilots [T3]

Description: A threat or failure scenario involving prompt injection in engineering copilots in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D2-CTL-05; D3-CTL-06; D4-CTL-06; D5-CTL-06

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-014 - Indirect prompt injection through service data [T3]

Description: A threat or failure scenario involving indirect prompt injection through service data in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D2-CTL-06; D3-CTL-07; D4-CTL-07; D6-CTL-01

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-015 - Excessive AI-agent authority [T3]

Description: A threat or failure scenario involving excessive ai-agent authority in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D3-CTL-01; D4-CTL-01; D5-CTL-01; D6-CTL-02

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-016 - Unsafe generated code [T3]

Description: A threat or failure scenario involving unsafe generated code in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D3-CTL-02; D4-CTL-02; D5-CTL-02; D6-CTL-03

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-017 - Hallucinated diagnostic guidance [T3]

Description: A threat or failure scenario involving hallucinated diagnostic guidance in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D3-CTL-03; D4-CTL-03; D5-CTL-03; D6-CTL-04

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-018 - Driver-monitoring evasion [T3]

Description: A threat or failure scenario involving driver-monitoring evasion in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D3-CTL-04; D4-CTL-04; D5-CTL-04; D6-CTL-05

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-019 - Vehicle-data leakage [T3]

Description: A threat or failure scenario involving vehicle-data leakage in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D3-CTL-05; D4-CTL-05; D5-CTL-05; D6-CTL-06

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-020 - Location-data misuse [T3]

Description: A threat or failure scenario involving location-data misuse in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D3-CTL-06; D4-CTL-06; D5-CTL-06; D6-CTL-07

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-021 - Fleet-wide correlated model failure [T3]

Description: A threat or failure scenario involving fleet-wide correlated model failure in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D3-CTL-07; D4-CTL-07; D6-CTL-01; D7-CTL-H01

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-022 - Automation bias [T3]

Description: A threat or failure scenario involving automation bias in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D4-CTL-01; D5-CTL-01; D6-CTL-02; D7-CTL-H02

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-023 - Inadequate fallback behaviour [T3]

Description: A threat or failure scenario involving inadequate fallback behaviour in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D4-CTL-02; D5-CTL-02; D6-CTL-03; D7-CTL-H03

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-024 - Loss of observability [T3]

Description: A threat or failure scenario involving loss of observability in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D4-CTL-03; D5-CTL-03; D6-CTL-04; D7-CTL-H04

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-025 - Insecure OTA infrastructure [T3]

Description: A threat or failure scenario involving insecure ota infrastructure in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D4-CTL-04; D5-CTL-04; D6-CTL-05; D7-CTL-H05

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-026 - Signing-key compromise [T3]

Description: A threat or failure scenario involving signing-key compromise in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D4-CTL-05; D5-CTL-05; D6-CTL-06; D8-CTL-01

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-027 - Supplier compromise [T3]

Description: A threat or failure scenario involving supplier compromise in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D4-CTL-06; D5-CTL-06; D6-CTL-07; D8-CTL-02

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-028 - Shadow AI in engineering [T3]

Description: A threat or failure scenario involving shadow ai in engineering in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D4-CTL-07; D6-CTL-01; D7-CTL-H01; D8-CTL-03

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-029 - Manipulated test evidence [T3]

Description: A threat or failure scenario involving manipulated test evidence in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D5-CTL-01; D6-CTL-02; D7-CTL-H02; D8-CTL-04

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-030 - Post-release model drift [T3]

Description: A threat or failure scenario involving post-release model drift in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D5-CTL-02; D6-CTL-03; D7-CTL-H03; D8-CTL-05

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-031 - Environmental distribution shift [T3]

Description: A threat or failure scenario involving environmental distribution shift in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D5-CTL-03; D6-CTL-04; D7-CTL-H04; D9-CTL-01

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-032 - Rare-event failure [T3]

Description: A threat or failure scenario involving rare-event failure in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D5-CTL-04; D6-CTL-05; D7-CTL-H05; D9-CTL-02

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-033 - Unsafe AI-deterministic control interaction [T3]

Description: A threat or failure scenario involving unsafe ai-deterministic control interaction in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D5-CTL-05; D6-CTL-06; D8-CTL-01; D9-CTL-03

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

ATS-034 - Cross-discipline assurance failure [T3]

Description: A threat or failure scenario involving cross-discipline assurance failure in an automotive AI lifecycle or operating environment.

Initiating condition: External attacker, malicious insider, compromised supplier, defective process, environmental condition, or unintentional operator action.

Pathway: Compromise or degradation of data, model, software, interface, update, identity, monitoring, or governance controls.

Consequences: Human - Potential ranges from inconvenience or privacy harm to injury depending on vehicle authority and fallback capability.; Security - Loss of confidentiality, integrity, availability, authenticity, or accountability.; Operational/mission - Degraded vehicle, fleet, engineering, manufacturing, or service objective.; Financial - Investigation, remediation, recall, downtime, litigation, contractual, or regulatory cost.

Relevant controls: D5-CTL-06; D6-CTL-07; D8-CTL-02; D9-CTL-04

Observed status: Plausible scenario; not asserted as a confirmed field incident.

Evidence tier: T3 - analytical or research-supported

Limitations: Occurrence, exploitability, and consequence depend on implementation and operating context.

Annex D - Evidence catalogue

AEV-001 - AI system inventory

Purpose: Demonstrate the design, implementation, operation, or review of ai system inventory.

Owner: AI Governance Lead

Lifecycle: Concept; use-case approval; change control

Control linkage: D6-CTL-02; D6-CTL-06; D8-CTL-01; D8-CTL-02

Linkage justification: Establishes the governed population to which controls, assessments, exceptions and review obligations apply.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-002 - System boundary record

Purpose: Demonstrate the design, implementation, operation, or review of system boundary record.

Owner: Systems Engineering Lead

Lifecycle: System definition; architecture; integration

Control linkage: D3-CTL-01; D3-CTL-02; D3-CTL-04; D9-CTL-01; D9-CTL-02; D9-CTL-06

Linkage justification: Defines interfaces, authority, trust boundaries and physical-control influence needed to test control applicability.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-003 - Automotive applicability decision

Purpose: Demonstrate the design, implementation, operation, or review of automotive applicability decision.

Owner: AI Governance Lead

Lifecycle: Use-case approval; concept; management review

Control linkage: D6-CTL-01; D6-CTL-02; D8-CTL-01; D8-CTL-02; D9-CTL-01

Linkage justification: Records why each control applies, does not apply or requires enhanced treatment in the automotive context.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-004 - Architecture and interface description

Purpose: Demonstrate the design, implementation, operation, or review of architecture and interface description.

Owner: Systems Engineering Lead

Lifecycle: Architecture; design; integration

Control linkage: D2-CTL-05; D2-CTL-06; D3-CTL-01; D3-CTL-02; D4-CTL-05; D9-CTL-06

Linkage justification: Supports verification of interfaces, tool permissions, trust boundaries and command paths.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-005 - Threat analysis and risk assessment

Purpose: Demonstrate the design, implementation, operation, or review of threat analysis and risk assessment.

Owner: Automotive Cybersecurity Lead

Lifecycle: Concept; cybersecurity analysis; change assessment

Control linkage: D1-CTL-01; D1-CTL-03; D2-CTL-01; D2-CTL-02; D3-CTL-04; D4-CTL-05; D6-CTL-04; D9-CTL-04

Linkage justification: Documents attack paths, initiating conditions, consequence vectors, controls and residual risk.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-006 - Functional-safety coordination record

Purpose: Demonstrate the design, implementation, operation, or review of functional-safety coordination record.

Owner: Functional Safety Lead

Lifecycle: Safety analysis; verification; release; material change

Control linkage: D3-CTL-04; D6-CTL-01; D9-CTL-01; D9-CTL-02; D9-CTL-03; D9-CTL-06; D9-CTL-07

Linkage justification: Shows that AI-security changes and assumptions were reconciled with functional-safety artefacts and release decisions.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-007 - SOTIF coordination record

Purpose: Demonstrate the design, implementation, operation, or review of sotif coordination record.

Owner: SOTIF Lead

Lifecycle: SOTIF analysis; scenario validation; release; field monitoring

Control linkage: D1-CTL-03; D3-CTL-04; D6-CTL-01; D9-CTL-01; D9-CTL-02; D9-CTL-05

Linkage justification: Shows treatment of performance limitations, foreseeable misuse, triggering conditions and unknown unsafe scenarios.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-008 - Cybersecurity concept or case

Purpose: Demonstrate the design, implementation, operation, or review of cybersecurity concept or case.

Owner: Automotive Cybersecurity Lead

Lifecycle: Concept; architecture; cybersecurity validation

Control linkage: D2-CTL-01; D2-CTL-02; D2-CTL-04; D3-CTL-02; D4-CTL-05; D9-CTL-04; D9-CTL-06

Linkage justification: Provides the automotive cybersecurity rationale, control strategy and assurance argument.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-009 - Data provenance record

Purpose: Demonstrate the design, implementation, operation, or review of data provenance record.

Owner: Data Governance Lead

Lifecycle: Data acquisition; labelling; model development; retraining

Control linkage: D1-CTL-01; D1-CTL-04; D5-CTL-02; D5-CTL-05; D5-CTL-06

Linkage justification: Demonstrates origin, rights, transformations, quality controls, access and lineage of training and evaluation data.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-010 - Model card

Purpose: Demonstrate the design, implementation, operation, or review of model card.

Owner: AI Model Owner

Lifecycle: Model development; validation; release; material change

Control linkage: D1-CTL-03; D6-CTL-03; D8-CTL-03

Linkage justification: Records intended use, limitations, evaluation basis, dependencies, release status and known risks.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-011 - Model provenance and bill of materials

Purpose: Demonstrate the design, implementation, operation, or review of model provenance and bill of materials.

Owner: AI Security Lead

Lifecycle: Acquisition; model development; integration; release; update

Control linkage: D1-CTL-06; D1-CTL-07; D1-CTL-08; D1-CTL-09; D4-CTL-01; D4-CTL-02; D4-CTL-03; D4-CTL-07

Linkage justification: Establishes model/component identity, provenance, dependency, licence, integrity and approved release lineage.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-012 - Supplier assessment

Purpose: Demonstrate the design, implementation, operation, or review of supplier assessment.

Owner: Supplier Assurance Lead

Lifecycle: Supplier sourcing; onboarding; periodic review; change

Control linkage: D4-CTL-01; D4-CTL-02; D4-CTL-03; D4-CTL-05; D4-CTL-07; D6-CTL-06

Linkage justification: Evaluates supplier capability, evidence, incident obligations, support period, subcontractors and change controls.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-013 - Contractual security requirements

Purpose: Demonstrate the design, implementation, operation, or review of contractual security requirements.

Owner: Legal Counsel and Supplier Assurance Lead

Lifecycle: Contracting; sourcing; renewal; change

Control linkage: D4-CTL-05; D6-CTL-06; D8-CTL-01; D8-CTL-03; D8-CTL-05

Linkage justification: Creates enforceable obligations for security evidence, notification, audit, support, change and end-of-life.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-014 - Robustness test report

Purpose: Demonstrate the design, implementation, operation, or review of robustness test report.

Owner: AI Validation Lead

Lifecycle: Model verification; robustness testing; release; material change

Control linkage: D1-CTL-03; D1-CTL-05; D2-CTL-03; D2-CTL-04; D9-CTL-02; D9-CTL-05

Linkage justification: Demonstrates performance and failure behaviour across representative perturbations, environments and degradation conditions.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-015 - Adversarial test report

Purpose: Demonstrate the design, implementation, operation, or review of adversarial test report.

Owner: AI Security Lead

Lifecycle: Security validation; red teaming; release; material change

Control linkage: D1-CTL-01; D1-CTL-02; D2-CTL-01; D2-CTL-02; D2-CTL-03; D2-CTL-04; D2-CTL-05; D2-CTL-06; D9-CTL-04

Linkage justification: Records adversarial test scope, methods, results, exploitability, limitations, remediation and retest evidence.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-016 - Scenario-coverage report

Purpose: Demonstrate the design, implementation, operation, or review of scenario-coverage report.

Owner: SOTIF Lead and AI Validation Lead

Lifecycle: Scenario development; simulation; verification; field update

Control linkage: D1-CTL-03; D3-CTL-04; D9-CTL-01; D9-CTL-02; D9-CTL-05

Linkage justification: Shows coverage of operational, environmental, misuse, rare-event and safety-relevant scenarios.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-017 - Software verification report

Purpose: Demonstrate the design, implementation, operation, or review of software verification report.

Owner: Vehicle Software Lead

Lifecycle: Software verification; integration; production release; update

Control linkage: D1-CTL-06; D2-CTL-05; D3-CTL-01; D4-CTL-02; D4-CTL-07; D8-CTL-05; D9-CTL-06

Linkage justification: Demonstrates software and integration verification for the approved build, configuration and command paths.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-018 - OTA release and signing record

Purpose: Demonstrate the design, implementation, operation, or review of ota release and signing record.

Owner: Vehicle Software and OTA Lead

Lifecycle: Production release; OTA deployment; rollback; post-update review

Control linkage: D1-CTL-06; D1-CTL-07; D1-CTL-08; D1-CTL-09; D4-CTL-01; D4-CTL-02; D9-CTL-06

Linkage justification: Proves authorised signing, release approval, compatibility, deployment scope, rollback readiness and update traceability.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-019 - Monitoring configuration

Purpose: Demonstrate the design, implementation, operation, or review of monitoring configuration.

Owner: Vehicle Security Operations Lead

Lifecycle: Deployment; monitoring; incident detection; tuning

Control linkage: D1-CTL-03; D4-CTL-04; D6-CTL-02; D6-CTL-04; D6-CTL-07; D9-CTL-04; D9-CTL-05

Linkage justification: Defines monitored signals, detection logic, thresholds, escalation, retention and known visibility gaps.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-020 - Fleet anomaly report

Purpose: Demonstrate the design, implementation, operation, or review of fleet anomaly report.

Owner: Vehicle Security Operations Lead

Lifecycle: Fleet operations; monitoring; incident response; field action

Control linkage: D1-CTL-03; D4-CTL-04; D6-CTL-04; D6-CTL-07; D9-CTL-04; D9-CTL-05; D9-CTL-07

Linkage justification: Supports fleet-level correlation, common-mode detection, triage, containment and post-event analysis.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-021 - Incident record

Purpose: Demonstrate the design, implementation, operation, or review of incident record.

Owner: Incident Response Lead

Lifecycle: Incident response; product-safety escalation; recovery; lessons learned

Control linkage: D6-CTL-02; D6-CTL-04; D6-CTL-07; D7-CTL-H04; D9-CTL-04; D9-CTL-07

Linkage justification: Preserves event facts, decisions, communications, containment, recovery, safety escalation and corrective actions.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-022 - Exception and risk-acceptance record

Purpose: Demonstrate the design, implementation, operation, or review of exception and risk-acceptance record.

Owner: Risk Owner

Lifecycle: Risk treatment; exception approval; periodic review; expiry

Control linkage: D6-CTL-01; D6-CTL-02; D6-CTL-06; D6-CTL-07; D8-CTL-01; D9-CTL-01

Linkage justification: Records the accountable acceptance decision, compensating controls, monitoring, expiry and residual risk.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-023 - Training and competence record

Purpose: Demonstrate the design, implementation, operation, or review of training and competence record.

Owner: Competence and Training Lead

Lifecycle: Role assignment; onboarding; periodic training; change

Control linkage: D6-CTL-01; D6-CTL-04; D6-CTL-06; D7-CTL-H01; D7-CTL-H02; D7-CTL-H03; D7-CTL-H04

Linkage justification: Demonstrates that assigned personnel have role-relevant competence and current training.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-024 - Management review record

Purpose: Demonstrate the design, implementation, operation, or review of management review record.

Owner: Executive Risk Committee

Lifecycle: Periodic management review; major release; material incident

Control linkage: D6-CTL-02; D6-CTL-04; D6-CTL-06; D6-CTL-07; D8-CTL-01; D8-CTL-02

Linkage justification: Documents leadership review of performance, exceptions, incidents, resources, residual risk and corrective action.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

AEV-025 - Corrective-action record

Purpose: Demonstrate the design, implementation, operation, or review of corrective-action record.

Owner: Control Owner

Lifecycle: Remediation; verification; closure; continual improvement

Control linkage: D1-CTL-03; D4-CTL-02; D6-CTL-02; D6-CTL-04; D6-CTL-07; D8-CTL-02; D8-CTL-05; D9-CTL-07

Linkage justification: Links findings to accountable remediation, verification evidence, closure criteria and recurrence prevention.

Integrity / limitations: Version controlled, attributable, tamper-evident where material, and linked to approved scope. Possession of the artefact does not alone prove control effectiveness.

Annex E - Notably Absent register

- No claim of automatic legal compliance, homologation, type approval, functional-safety approval, SOTIF approval, cybersecurity approval, accredited certification, or prevention of harm.
- No verified field exploitation is asserted for illustrative threat scenarios unless separately supported by a primary source.
- No clause-level equivalence mapping is asserted where full standards text was not lawfully accessible and reviewed.
- No proprietary vehicle telemetry, client information, or internal incident dataset was used.
- No specialist review gate is represented as closed without documented findings and closure evidence.

Annex F - Open specialist review gates

The package retains 16 open gates: automotive cybersecurity engineering; functional safety; SOTIF; vehicle software and OTA; automated driving; product safety; homologation and type approval; automotive privacy; AI security; legal and regulatory; supplier assurance; assessment methodology; certification-scheme alignment; accessibility; technical editing; and final publication approval.

Limitations and publication status

This document is an informative sector implementation layer. It does not replace the authoritative GAISSE standard, applicable law, vehicle engineering, competent safety judgement, type approval, homologation, cybersecurity engineering, privacy review, or independent assessment. Publication status remains **Publication Candidate - Open Review Gates**.