GAIA-X 4 AGEDA: A Federated and Secure Cloud Framework with AI-Driven Traffic Compliance in Connected Vehicles

Achim Rettberg, Narmada Ambigapathy, Fatima Idrees, Ali Alhalabi, Katrin Glöwing, Mehdi Azarafza
Hamm-Lippstadt University of Applied Sciences, Lippstadt, Germany
Emails: {achim.rettberg, narmada.ambigapathy, fatima.idrees, ali.alhalabi, katrin.gloewing, mehdi.azarafza}@hshl.de

***Abstract*—Connected and autonomous vehicles increasingly rely on secure cloud infrastructures to exchange perception, safety, and compliance information. Yet, the lack of standardized mechanisms for trustworthy data federation limits interoperability and regulatory validation across heterogeneous vehicle fleets. This paper presents GAIA-X 4 AGEDA, a federated and secure cloud framework that combines dynamic edge orchestration with AI-driven traffic compliance monitoring for connected vehicles. The architecture leverages the GAIA-X Trust Framework for credential-based authentication and sovereign data exchange, while integrating Orchestration for mixed-criticality workload management across a four-zone vehicle architecture. An edge-AI compliance module validates real-time adherence to legal road traffic regulations, with results securely transmitted through GAIA-X Data Spaces to enable cooperative awareness among connected vehicles. Using a CARLA-based proof-of-concept [1] deployed on Raspberry Pi 5 and Jetson Nano P3541 platforms, we demonstrate how compliance violations detected at the edge are authenticated, federated, and shared with Traffic Awareness Services (TAS) while preserving data sovereignty. The framework achieves V2X compliance alerts and this work establishes a practical pathway toward dependable, regulation-aware, and federated vehicle ecosystems aligned with European data sovereignty principles, enabling secure service deployment throughout the vehicle lifecycle.**

***Index Terms*—GAIA-X 4 AGEDA, federated data spaces, connected vehicles, traffic compliance, edge orchestration, AI-driven monitoring, cloud federation, software-defined vehicles**

I. Introduction

Modern vehicles have evolved into software-defined, interconnected entities capable of executing autonomous decisions and real-time communication with surrounding infrastructure. This transformation introduces critical challenges in **data security, regulatory compliance, and trust management**. Autonomous vehicles must continuously exchange safety-critical information with other vehicles and cloud services which requires trusted mechanisms for data authentication, validation, and sovereignty aligned with European Union (EU) digital-sovereignty principles. Figure 1 highlights how trust, interoperability, and data sovereignty are embedded within the proposed architecture.

The GAIA-X 4 AGEDA project, part of the GAIA-X for Future Mobility initiative [2], addresses these challenges through a federated, secure, and interoperable cloud framework. GAIA-X ensures all participants (vehicles, OEMs, service providers) adhere to a unified trust model, leveraging credential verification, digital identity, and compliance auditing [3]. Simultaneously, AI-driven systems based on Large Language Models (LLMs) are emerging as novel tools for traffic management and monitoring, yet often operate in isolation from trusted cloud infrastructures.

This paper was partly funded by German Federal Ministry for Economic Affairs and Energy (BMWE 19S22004O) within the GAIA-X 4 AGEDA project. The sole responsibility for the content lies with the authors.

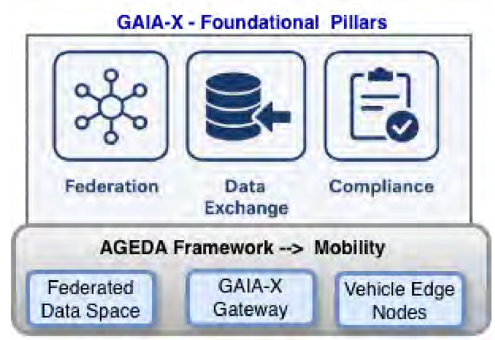

Fig. 1: Conceptual alignment between GAIA-X foundational pillars and the AGEDA mobility framework.

This paper bridges that gap by presenting a GAIA-X 4 AGEDA-based architecture integrating AI-driven traffic-compliance verification within a federated data space.

Main contributions of this work include:

- A **GAIA-X-compliant federated architecture** enabling vehicles to function as authenticated edge participants.
- **Integration of AI-driven monitoring/ compliance module** with GAIA-X Data Spaces for secure, sovereign data exchange.
- A **proof-of-concept demonstration** using CARLA-based simulation integrated with Raspberry Pi 5, and Jetson Nano platforms demonstrating road regulation monitoring aided with **credential-based authentication**.

II. Background and Related Work

A. Federated Vehicle Architectures and Trust Frameworks

Software-Defined Vehicles (SDVs) represent the convergence of embedded systems and cloud technologies, enabling continuous updates and service orchestration. The AGEDA framework integrates Eclipse Ankaios for orchestration engine across a four-zone vehicle architecture—perception, decision, control, and communication enabling dynamic workload migration between edge nodes and cloud instances [4].

GAIA-X establishes a federated ecosystem where data providers and consumers retain sovereignty over their assets [3]. It's Trust Framework governs credential issuance and compliance verification through components such as the GAIA-X Registry, Authentication Service, and Digital Clearing House. Within mobility, GAIA-X enables vehicles to act as trusted participants under self-sovereign identity (SSI) principles, while the AGEDA project adapts these mechanisms for secure interoperability among OEMs, service providers, and smart-city infrastructure.

B. AI-Driven Traffic Compliance and Edge Intelligence

Recent advances highlight the potential of Large Language Models (LLMs) in traffic management due to their context aware reasoning capabilities. Studies have applied LLMs to assess rule violations and explain driving behavior. However, their cloud dependency and high inference latency hinder real-time use. To mitigate this, edge deployment of optimized models through quantization and hybrid edge–cloud orchestration has gained traction [5]. Emerging frameworks now integrate semantic edge intelligence with trusted data sharing; for instance, Onsu et al. [6] introduced a semantic edge–cloud system using Vision Transformers (ViTs) and LLMs that compress visual features into compact embeddings, achieving a 99.9% reduction in transmission size while maintaining reasoning accuracy. Such developments lay the foundation for AI-driven traffic compliance frameworks [7] that couple decentralized edge-deployed context aware reasoning with data exchange to ensure regulatory adherence and cooperative awareness. Since, these AI modules are typically constrained to local inference without secure federation there lies gap in approaches that unifies **AI-based compliance reasoning** with a **trusted, credential-based data-sharing layer**. GAIA-X 4 AGEDA addresses this gap by embedding compliance verification within a federated trust and governance architecture.

III. Architecture

A. System Overview

The GAIA-X 4 AGEDA architecture (Figure 2) comprises four principal layers:

Vehicle Edge Layer hosts in-vehicle computation including control, perception, V2X communication, and AD Stack, handling sensor fusion and inference locally for low-latency responses.

AGEDA Framework (Middleware) manages workload orchestration and data flow through the Orchestration Engine, MQTT/gRPC Gateway, GAIAX-Connector for credential-based access, and Platform Services (Kuksa Databroker) for real-time data availability.

Application Layer enforces trust through MoveID [8] (SSI-based authentication) and Credential Manager, while providing Fleet Management Dashboard, Regulatory Compliance Portal, and Traffic Awareness Service.

GAIA-X Data Space (Federated Cloud) includes GAIA-X Registry, Digital Clearing House (GX-DCH), Federated Catalogue, and Data Space for credential validation and sovereign storage. This design enables bidirectional trust flow from federated cloud governance to edge execution.

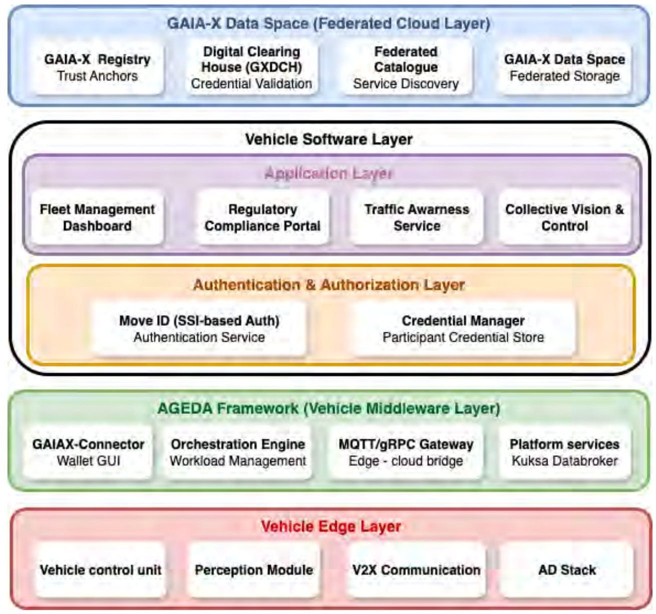

Fig. 2: GAIA-X 4 AGEDA federated vehicle architecture showing interactions between edge, middleware, authentication, and federated cloud layers.

B. Data Flow and Interaction Pipeline

The data flow illustrated in Figure 3 depicts the complete lifecycle of an AI-driven traffic compliance event within the GAIA-X 4 AGEDA framework. The pipeline integrates contextual reasoning, credential-based authentication, and federated dissemination to ensure trustworthy, low-latency information sharing among connected vehicles.

1) Compliance Detection: Edge AI agents process multi-modal perception data to classify behavior according to regulations, producing JSON-LD objects with regulation reference, compliance status, confidence score, and timestamp.

2) Messaging Structure: Events are serialized to GAIA-X-compliant JSON-LD format enriched with vehicle identification (vehicleID), Global Positioning System (GPS), and timestamp for semantic interoperability.

3) Secure Transmission: Encrypted messages transmit via TLS-encrypted MQTT/gRPC to the GAIA-X Gateway, ensuring confidentiality and authenticity.

4) Credential Verification: GAIA-X Authentication Service verifies sender credentials using MoveID and GX-DCH. Only authorized vehicles publish compliance data.

5) Federated Distribution: Verified events publish to GAIA-X Data Space, accessible to TAS, nearby vehicles, and stakeholders. Policy-controlled access and audit logging enable V2X alerts for cooperative awareness.

C. Security, Interoperability, and Governance

The framework designed to adhere to EU-standard principles for **trust, transparency, and interoperability**:

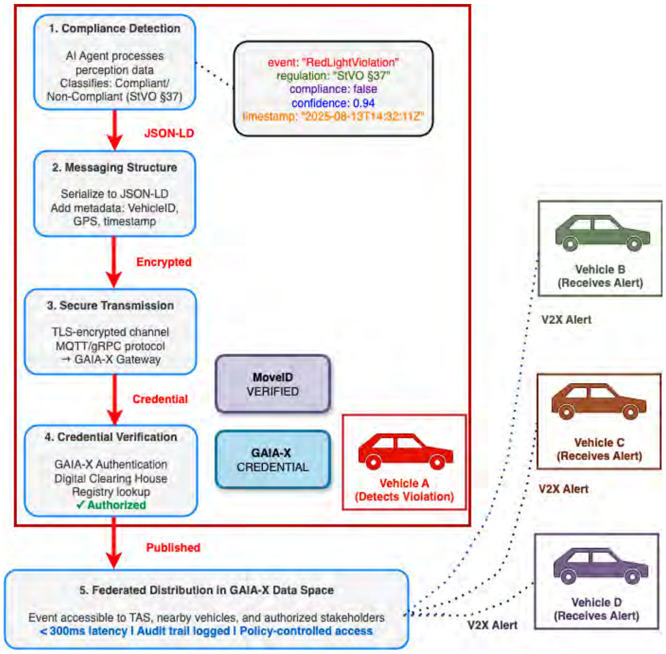

Fig. 3: Federated V2X compliance data exchange pipeline within the GAIA-X 4 AGEDA framework.

The framework ensures:

- **Credential-based Access Control:** Using moveID and GAIA-X credentials
- **End-to-End Encryption:** TLS and digital signatures protect all exchanges
- **Semantic Interoperability:** JSON-LD schemas provide machine-readable context
- **Auditability:** Transaction logging enables regulatory traceability

This ensures that compliance messages are verifiable, tamper-proof, and interoperable, supporting dependable behavior in multi-vehicle environments.

IV. Proof-of-Concept Demonstration

A. Simulation Setup

To validate the proposed architecture, a digital-twin testbed was implemented using the CARLA simulator(version 0.9.15) integrated with the GAIA-X framework. The simulation environment reproduced urban intersections, multilane roads, and pedestrian crossings consistent with German Road Traffic Regulations (StVO [9]), enabling evaluation of cooperative compliance awareness in connected vehicle platoons.

Edge inference was executed on two embedded platforms Raspberry Pi 5 and Jetson Nano P3541, representing resource-constrained vehicle nodes. Each device hosted a fine-tuned LLM based compliance agent. The fine-tuned models included four large models (GPT-2 [10], Neo [11], OPT [12], BLOOM [13]) and two lightweight models (DistilGPT2 [14], FLAN-T5-small [15]). The compliance agents processed textual summaries generated from CARLA's perception data streams (vehicle state, throttle, steering, and environment context) and produced structured outputs identifying Compliant or Non-Compliant behavior. A lightweight MQTT broker enabled asynchronous message exchange between CARLA and the edge nodes, while the GAIA-X Trust Layer managed credential verification and secure routing toward the GAIA-X Gateway. Each inference result was serialized in JSON-LD, encrypted, and transmitted to the federated cloud for authentication before publication within the GAIA-X Data Space. Although end-to-end system latency was not formally benchmarked across all scenarios, component-level measurements indicated aggregate delays suitable for non-critical cooperative awareness applications.

B. Demonstration Scenarios

Three scenarios validated federated compliance capabilities across connected vehicles navigating urban environments.

Scenario 1 – Stop Sign Recognition and Coordinated Response (Figure 4): A lead vehicle approaching a stop sign performs a complete halt per StVO § 8 and § 41, Sign 206, which the on-board edge agent classifies as Compliant. The signed JSON-LD event is sent via the GAIA-X Gateway; after credential verification and Digital Clearing House approval, it is published to the Data Space. Following vehicles receive the federated notice and pre-emptively adapt, demonstrating synchronized, rule-aware behavior across authenticated participants.

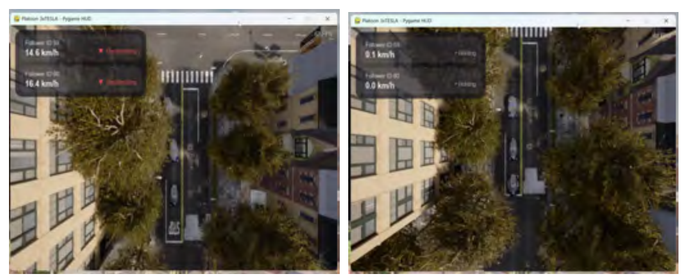

(a) Scenario 1 - a (b) Scenario 1 - b

Fig. 4: (a) Stop-sign compliance detected by the lead vehicle; (b) synchronized halt of following vehicles after federated event dissemination through the GAIA-X Data Space (StVO §8, §41 Sign 206).

s

Fig. 5: Scenario 2: Traffic Signal Violation: Speed adjustment as per StVO Section 37

Scenario 2 – Red Light Violation Detection and Alert Propagation (Figure 5): When a vehicle fails to stop at a red signal, the compliance agent classifies the behavior as Non-Compliant with StVO §37. The violation event is transmitted for credential validation, then broadcast via TAS to approaching vehicles, enabling hazard anticipation and trajectory modification.

Scenario 3 – Multi-Vehicle Platooning with Cooperative Speed Adjustment (Figure 6): A lead vehicle navigating moderate traffic publishes continuous compliance updates regarding speed limits (StVO §3(1)). Following vehicles in the platoon synchronize their velocity profiles based on federated speed adjustment messages, maintaining safe inter-vehicle spacing while collectively adhering to regulatory constraints. This demonstrates sustained cooperative awareness in dynamic traffic conditions.

Fig. 6: Scenario 3: Platooning while adjusting speed as per StVO Section 3(1): Speed Adjustment while navigating moderate traffic

V. Discussion

The GAIA-X 4 AGEDA framework bridges AI-driven compliance with federated trust infrastructure, addressing critical gaps in connected vehicle ecosystems. By federating high-level compliance interpretations rather than raw telemetry, the architecture reduces bandwidth overhead while ensuring verifiable provenance through credential-based validation, essential for liability determination and regulatory acceptance. The demonstrated cooperative awareness mechanism shifts traffic enforcement from isolated checkpoints to continuous validation, enabling proactive compliance assistance while preserving data sovereignty through policy-controlled Data Spaces accessible to multi-stakeholder participants (OEMs, municipalities, authorities) without central ownership.

Key limitations include simulation-based validation requiring real-world testing, LLM brittleness to out-of-distribution scenarios necessitating hybrid formal verification approaches, and scalability challenges for million-vehicle deployments demanding hierarchical federation architectures. Future directions will include federated learning for decentralized model improvement, production grade Eclipse Ankaios integration, and cross-domain extensions to smart city and logistics infrastructures.

VI. Conclusion

This paper presented GAIA-X 4 AGEDA, a federated cloud framework integrating AI-driven traffic compliance for connected vehicles through GAIA-X Trust Framework components and edge-deployed LLM agents. The CARLA-based proof-of-concept demostrated complete information flow from edge detection through credential verification to federated dissemination across three scenarios demonstrating cooperative awareness, traffic violation detection, and platooning coordination. By merging edge intelligence with federated trust, this work introduces a practical pathway toward regulation-aware vehicle ecosystems aligned with European data sovereignty principles, enabling secure, dependable mobility services throughout the vehicle lifecycle.

References

[1] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, CARLA: An Open Urban Driving Simulator, arXiv preprint arXiv:1711.03938, 2017. [Online]. Available: https://arxiv.org/abs/1711.03938#

[2] GAIA-X 4 Future Mobility, "GAIA-X 4 AGEDA Project," Accessed: Oct. 10, 2025. [Online]. Available: https://www.gaia-x4futuremobility.de/en/projects/ageda

[3] GAIA-X AISBL, "GAIA-X Framework," Accessed: Oct. 10, 2025. [Online]. Available: https://gaia-x.eu/gaia-x-framework/

[4] N. Ambigapathy, C. Steinmetz, L. Stahlbock, K. Rusev, A. Kos, S. A. Böttigheimer, and A. Rettberg, "GAIA-X 4 AGEDA Powered Software-Defined Vehicles: Hybrid Failure-Aware Orchestration for Adaptive Mixed-Criticality Workload Management," in *Proc. (IAVVC)*, pp. 1–6. [Online]. Available: https://doi.org/10.1109/IAVVC61942.2025.11219638

[5] S. Semerikov, T. Vakaliuk, O. Kanevska, M. Moiseienko, I. Donchev, and A. Kolhatin, "LLM on the edge: the new frontier," in *Proc. CEUR Workshop Vol. 3943*, Mar. 2025.

[6] M. A. Onsu, P. Lohan, B. Kantarci, A. Syed, M. Andrews, and S. Kennedy, "Semantic edge–cloud communication for real-time urban traffic surveillance with ViT and LLMs over mobile networks," *arXiv preprint arXiv:2509.21259*, 2025. [Online]. Available: https://arxiv.org/abs/2509.21259

[7] F. Idrees, N. Ambigapathy, P. Adelt, and A. Rettberg, "AI-Driven Traffic Compliance in Autonomous Vehicles: A Scalable Framework with Large Language Models on Edge Platforms," in *Proc. (IAVVC)*, pp. 1–6. [Online]. Available: https://doi.org/10.1109/IAVVC61942.2025.11219567

[8] GAIA-X - "MoveID Project," Accessed: Oct. 10, 2025. [Online]. Available: https://www.gaia-x4futuremobility.de/en/projects/moveid

[9] German Road Traffic Regulations (StVO), Bundesministerium für Digitales und Verkehr.

[10] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," in Proceedings, 2019.

[11] L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy, "The Pile: An 800GB dataset of diverse text for language modeling," CoRR, vol. abs/2101.00027, 2021.

[12] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, T. Mihaylov, M. Ott, S. Shleifer, K. Shuster, D. Simig, P. S. Koura, A. Sridhar, T. Wang, and L. Zettlemoyer, "OPT: Open pre-trained transformer language models," arXiv, 2022.

[13] BigScience Workshop, "BLOOM: A 176B-parameter open-access multilingual language model," arXiv e-prints, no. arXiv:2211.05100, 2022.

[14] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilGPT-2: a distilled, faster, and lighter version of GPT-2," Hugging Face, 2019. [Online]. Available: https://huggingface.co/distilbert/distilgpt2

[15] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, E. Li, X. Wang, M. Dehghani, S. Brahma, A. Webson, S. Gu, Z. Dai, M. Suzgun, X. Chen, A. Chowdhery, J. Devlin, A. Roberts, D. Zhou, Q. Le, J. Wei, "Scaling Instruction-Finetuned Language Models," arXiv preprint arXiv:2210.11416, 2022. [Online]. Available: https://arxiv.org/abs/2210.11416