AECC

AUTOMOTIVE EDGE
COMPUTING CONSORTIUM

Data-First Architecture for Data-Driven Automotive Service Development

Version 1.0 • April 15, 2026

Copyright © Automotive Edge Computing Consortium 2026

Table of Contents

Foreword **3**
1 Executive Summary **4**
2 Terminology **5**
3 Introduction **7**
3.1 Motivation 7
3.2 Key Issues of the Data-Driven Development Platform 8
3.3 Design Principles 9
3.4 Diversifying Communication Paths 9
4 Service Scenarios and Requirements **10**
4.1 AD / ADAS Model Training 10
4.1.1 Service Scenario Description 10
4.1.2 Uplink Data Characteristics and Requirements 10
4.2 Service Monitoring 12
4.2.1 Service Scenario Description 12
4.2.2 Uplink Data Characteristics and Requirements 13
4.2.3 Downlink Data Characteristics and Requirements 13
5 High-Level Network Architecture **14**
5.1 Layered Architecture 14
5.2 Data Transfer between Layers 15
5.3 Multipath Communications 16
5.3.1 Overview 16
5.3.2 Requirement Tags 17
5.3.3 Data Scheduler 18
5.3.4 Multipath Controller 18
5.3.5 Optimization of Individual Networks and Data Connections 18
6 Design Considerations **19**
6.1 Radio Access to Wi-Fi Access Points 19
6.1.1 Initial Link Setup Latency 19
6.1.2 Coverage Extension by Mesh Networking 20
6.2 Wi-Fi as an Enabler of Inter-Vehicle Data Transfer 20
6.2.1 Wi-Fi-based Peer-to-Peer Communications 20
6.2.2 Authentication for Inter-Vehicle Data Transfer 21
7 Conclusion and Next Steps **23**
8 References **24**

Foreword

"Data-First Architecture for Data-Driven Automotive Service Development" was prepared by the AECC's Data First Special Interest Group (SIG Data First).

The purpose of this document is to provide AECC's vision on scalable frameworks for collection, processing, and distribution of automotive big data, which constitute the foundation of automotive edge computing. In addition to the high-level architecture and example service scenarios, it provides a discussion of technical requirements, reference deployment patterns, and challenges to be addressed. This document is positioned as a supplement to the AECC's "General Principles and Vision" white paper [1].

AECC is a not-for-profit association of industry members, dedicated to promoting edge computing and communications technology in the connected vehicles ecosystem. For more information about AECC and its publications, visit https://aecc.org/.

1 Executive Summary

The Paradigm Shift to SDVs and the Data Explosion

The automotive industry is rapidly transitioning from a hardware-centric model to one defined by Software-Defined Vehicles (SDVs). Vehicles are no longer just modes of transport but intelligent edge devices tightly integrated with cloud environments. This evolution, driven by Over-The-Air (OTA) updates for continuous feature enhancement and the development of advanced AI models for ADAS (Advanced Driver Assistance Systems) and automated driving, has led to an exponential increase in the volume of automotive sensor and log data required for model training and validation. Traditional Internet-of-Things (IoT) architectures that rely heavily on cellular networks and public cloud infrastructure are struggling to cope with this massive scale, along with the associated costs and bandwidth bottlenecks [2].

Introducing the Data-First Architecture

To address these challenges, this white paper proposes the "Data-First Architecture," a communication and computation platform specifically designed for the scalable collection, processing, and distribution of automotive big data. The core principle is to position data at the center of the ecosystem and to utilize a highly distributed, tiered computing model. This architecture comprises three key layers:

- **Peer-to-Peer Network Layer:** Leverages opportunistic inter-vehicle data transfer to enable decentralized data aggregation and distribution among neighboring vehicles.
- **Edge Network Layer:** Incorporates Edge Data Centers and short-range wireless communication infrastructure (e.g., Wi-Fi access points) to opportunistically offload traffic and reduce the load on cellular networks.
- **Mobile and Cloud Network Layer:** Serves as the global layer for centralized data management, aggregation, and lifecycle control via cellular networks and cloud computing infrastructure.

Key Mechanisms: Multipath Communication and Delay-Tolerance

A crucial element of this architecture is Multipath Communication, which integrates diverse network paths (e.g., Cellular and Wi-Fi). The system intelligently classifies data based on its utility and delivery deadline (e.g., time-sensitive alerts vs. delay-tolerant training datasets). For large-volume, delay-tolerant traffic—such as raw sensor data for AD/ADAS model training or comprehensive service logs—a Data Scheduler temporarily stores the data locally and waits for more cost-effective offloading opportunities via Wi-Fi access points or inter-vehicle communication links, sending data over cellular communications only when deadlines approach or local connectivity is unavailable. This approach ensures service continuity by balancing network load and maximizing offload efficiency.

Scalability and Sustainability

The Data-First Architecture supports demanding use cases, including AD/ADAS model training (requiring tens of GB/day per vehicle) and service monitoring (requiring robust log collection and OTA updates). Furthermore, the architecture is designed for sustainability by reducing unnecessary data uploads and processing, promoting energy efficiency. Key technical design considerations include optimizing Wi-Fi radio access with features like FILS (Fast Initial Link Setup) and enhancing security for inter-vehicle communications. AECC aims to leverage this architecture to overcome current infrastructure constraints and accelerate the data-driven development of advanced, sustainable automotive services.

2 Terminology

Term	Definition
Confidence	The confidence level of an inference model or algorithm's estimate indicates the data's impact on downstream model training tasks.
Data Ferry	A service concept where vehicles collect data from other vehicles and IoT sensors along their travel routes, uploading it later via network infrastructure along roads and/or parking areas.
Data Scheduler	A component that evaluates data delivery deadlines and temporarily stores data locally to wait for cost-effective offloading opportunities (like Wi-Fi or inter-vehicle data transfer) or schedules batch transmissions over cellular networks.
Data-First Architecture	A communication and computation platform designed for the scalable collection, processing, and distribution of automotive big data, positioning data at the center of the ecosystem and utilizing a distributed, tiered computing model.
Edge Network Layer	The second layer of the architecture incorporates short-range wireless communication infrastructure (e.g., Wi-Fi access points) and Edge Data Centers to opportunistically offload traffic from cellular networks.
Fast Initial Link Setup (FILS)	A feature defined in the IEEE 802.11ai amendment that significantly decreases the number of control messages required for Wireless LAN authentication and performs IP address assignment during the authentication process to shorten link setup time.
Informativeness	A metric that measures how much a data sample reduces model uncertainty or improves decision boundaries when added to a training set.
Mobile and Cloud Network Layer	The global layer consists of cellular networks and cloud computing infrastructure responsible for centralized data management, aggregation, and lifecycle control.
Multipath Communication Gateway	A gateway that manages transmissions from vehicles, comprising the Data Scheduler and the Multipath Controller to handle data tagged with specific QoS requirements.
Multipath Controller	A component that selects the network interface(s) (e.g., cellular, Wi-Fi, inter-vehicle data transfer) for transmission based on requirement tags to ensure compliance with Quality of Service (QoS) parameters.
Peer-to-Peer Network Layer	The first layer of the architecture leverages inter-vehicle data transfer to enable decentralized data aggregation and distribution among neighboring vehicles.
Penalty	A metric indicating the operational impact associated with data, reflecting how critical the underlying event is and the urgency of propagating it to upper layers to avoid safety or reliability degradation.
Relevance	A metric indicating which service domains the data contributes to (such as AI training, monitoring, or mapping), representing the semantic value of the data within the vehicle service system.

Term	Definition
Requirement Tags	QoS tags (such as Deadline, Priority, QoS Class, and Sensitivity Class) attached to data by an application to define how the traffic should be handled by the Multipath Communication Gateway.
Software-Defined Vehicles (SDVs)	Vehicles that function as intelligent edge devices tightly integrated with cloud environments, where functions are defined and managed through software that can be updated Over-The-Air (OTA).
Token-based Decentralized Authentication	A security mechanism where a trusted certificate authority issues a digitally signed token to a vehicle, allowing it to authenticate locally with peers for inter-vehicle data transfer without contacting a remote server.
Uniqueness	A metric for data that captures rare events or scenarios not yet represented in the collected dataset, which may have a higher impact on development.
Wi-Fi Aware ™	A Wi-Fi Alliance specification for peer-to-peer communication that allows stations to establish direct one-to-one links without centralized coordinators like APs or Group Owners.
Wi-Fi Direct ®	A Wi-Fi Alliance specification enabling peer-to-peer communications in which certain stations serve as Group Owners (GOs) to function as temporary access points.
NAN Data Paths (NDP)	Direct one-to-one communication links are established between stations in a Wi-Fi Aware network without requiring central nodes, such as access points.
Delay-tolerant Data	Automotive data for which delivery timing is not critical and can tolerate significant transmission delays without affecting system functionality or safety. Such data may be delivered over extended timeframes, ranging from several hours to several months, and is typically associated with non-real-time functions such as diagnostics, usage statistics, or AI model training.
Time-sensitive Data	Automotive data that requires delivery within a strictly bounded time interval to ensure correct operation, performance, or safety. The acceptable latency for this data typically ranges from milliseconds to minutes, depending on the application, and is characteristic of real-time or near-real-time functions such as safety alerts and digital twin information for road traffic.

Note: Wi-Fi®, Wi-Fi Aware™, and Wi-Fi Direct® are trademarks or registered trademarks of Wi-Fi Alliance.

3 Introduction

3.1 Motivation

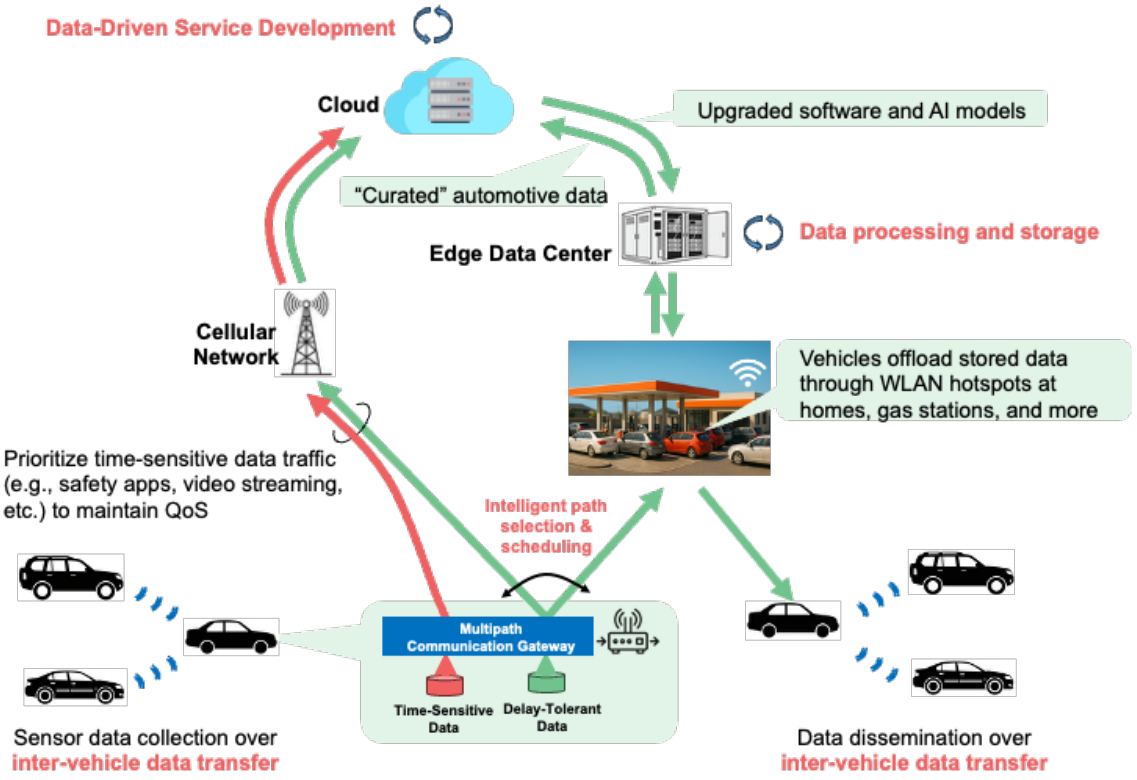

Figure 1: Overview of Data First Architecture

The rise of Software-Defined Vehicles (SDVs) has fundamentally transformed the delivery of automotive services. An increasing proportion of vehicle functions are now defined and managed through software, which can be hosted locally or remotely via cloud and edge computing infrastructure. Over-the-Air (OTA) software update capabilities enable continuous enhancements and feature expansion, allowing vehicles to evolve well beyond their initial release.

Automotive software development is becoming progressively more data-driven, with service logs and sensor data collected from vehicles and analyzed to assess the performance of new features and identify opportunities for improvement. The growing integration of AI-powered components further underscores the importance of this data-centric approach. AI has long been a cornerstone of advanced driver-assistance and automated driving systems, and its role is now extending into the infotainment domain, enabling intelligent voice assistants and delivering personalized travel recommendations. Training and optimizing these AI models require vast amounts of automotive data, often orders of magnitude more than traditional datasets. Consequently, they place far greater demands on data communication and processing than conventional Internet-of-Things (IoT) architectures, which have largely depended on cellular networks and public cloud infrastructure. Sustaining the pace of data-driven service innovation will require a scalable framework that efficiently collects and processes automotive big data.

In this paper, we introduce the **Data-First Architecture**, a communication and computation platform specifically designed to enable data-driven development of automotive services, with data positioned at the core of the ecosystem. Within this architecture, data is characterized by its large volume and tolerance to delays in collection and processing. For instance, in AI model training, massive datasets collected over defined time windows are

processed in batches to train and validate model parameters. The processing cadence may range from several days to several months, depending on development schedules. In such cases, vehicle-generated data does not need to be transmitted to data centers immediately; it is sufficient as long as the required data arrives before the next scheduled batch processing cycle.

However, wide-area networks such as cellular systems are not optimized for handling large-volume, delay-tolerant traffic, underscoring the need for an integrated solution that effectively combines different network types to meet diverse performance requirements. To address this gap, it is essential to leverage not only wide-area networks but also locally available communication resources. Inter-vehicle data transfer and localized systems enable data exchange and coordination within close proximity, allowing data to be aggregated, filtered, or propagated without immediately relying on centralized infrastructure. Such local-first interactions contribute to more efficient use of energy and network resources, while forming a cooperative domain in which vehicles collectively support data distribution and service continuity. From a system perspective, these approaches emphasize lightweight, long-lifetime connectivity mechanisms that can operate persistently and unobtrusively, yet meaningfully improve the overall efficiency and scalability of Software-Defined Vehicle (SDV) systems and their integration into broader societal infrastructures. By combining diverse communication paths spanning cellular, Wi-Fi, and inter-vehicle data transfer with a hierarchical edge computing model that integrates vehicle-, edge-, and cloud-based data processing, the Data First Architecture seeks to transcend the limitations of conventional IoT data platforms and accelerate the data-driven development of advanced automotive services.

3.2 Key Issues of the Data-Driven Development Platform

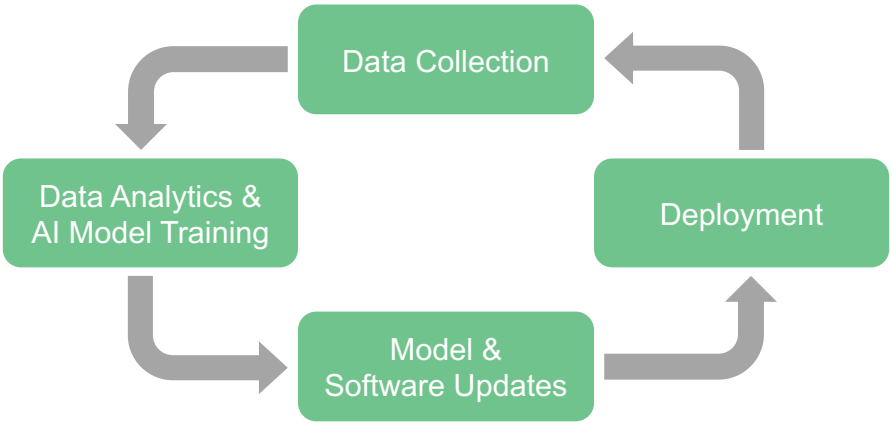

Figure 2: Iterative cycle of data-driven service development

The data-driven development process typically follows an iterative cycle consisting of the following steps:

(1) **Data collection**: Vehicles upload data to cloud computing infrastructure.

(2) **Data analytics and AI model training**: Collected data is analysed to train and refine AI models.

(3) **Model and software updates**: Upgraded AI models and software components are generated.

(4) **Deployment**: The updated content is distributed back to vehicles.

The data traffic involved in steps (1) and (4) differs significantly from that in service delivery platforms, which are optimized for providing services to end users. Data volumes are often orders of magnitude larger, sometimes exceeding tens of gigabytes per day, depending on the use case. Unlike service delivery systems, these processes typically have more relaxed latency requirements and can tolerate delays in data collection ranging from several days to even months. Wide-area network infrastructures, such as cellular networks, are not specifically optimized for this type of delay-tolerant, large-volume traffic. Although it is technically feasible to

upload such data over wide-area networks, doing so consumes substantial bandwidth and may degrade the performance of real-time services.

Furthermore, computing resources in cloud infrastructure are not unlimited. Therefore, steps (2) and (3) should not be entirely offloaded to the cloud. Instead, these functions should be strategically distributed across vehicle systems, edge data centers, and cloud platforms to balance performance, scalability, and resource efficiency. These observations highlight the need for a scalable data communication and processing platform capable of generating, collecting, storing, utilizing, and deleting automotive data in an *on-demand*, *flexible,* and *efficient* manner.

3.3 Design Principles

We adhere to the following design principles in the development of Data First Architecture:

- **Coordinated communications**: A vehicle consists of numerous software components. Allowing each component to communicate independently with remote servers can easily lead to network congestion. Therefore, communications within the vehicle system should be coordinated to optimize the number of active sessions and manage data transmission schedules effectively.
- **Scalable service monitoring**: Use cases that require continuous monitoring of service logs and anomaly detection should minimize system load by employing batched and/or scheduled data transmission.
- **Multipath communications**: Distributing and collecting large volumes of data across a fleet of vehicles can cause unexpected network loads if not properly managed. The data-driven development platform should therefore be able to balance network load across multiple communication paths and technologies to maintain service continuity.

3.4 Diversifying Communication Paths

Mobile network operators have made substantial capital investments to enhance base stations and core network infrastructure in support of advanced 5G mobile services. Despite these significant expenditures, it remains challenging to completely prevent large-scale network failures, which impose considerable burdens on both network operators and service providers that rely on these networks. Consequently, industry trends are shifting away from hosting users exclusively within a single operator’s domain toward a more collaborative model that shares resources and responsibilities across multiple operators offering heterogeneous network services—such as cellular, Wi-Fi, and non-terrestrial networks (NTNs). The hybrid use of diverse network types from multiple operators, a core driver of the Data First architecture, aligns well with the ongoing transformation of network operations and business models within the industry.

Another promising yet underutilized mode of communication is inter-vehicle data transfer. Through inter-vehicle communication links, vehicles can directly exchange information with nearby peers, enabling decentralized data distribution and aggregation. Although extensive research efforts have been devoted to this area for more than a decade, large-scale deployment remains at an early stage. To fully realize the scalability of data collection and dissemination networks, mechanisms for device-to-device cooperation must continue to evolve, with relevant functions integrated into wireless communication standards and supported by off-the-shelf communication hardware.

4 Service Scenarios and Requirements

The automotive industry is undergoing a paradigm shift from hardware-centric development to Software-Defined Vehicles (SDVs). In this new ecosystem, vehicles function as intelligent edge devices tightly integrated with cloud environments. The core of the SDV concept lies in establishing a continuous "DevOps loop": collecting vehicle data, training AI models in the cloud, and deploying improvements via Over-The-Air (OTA) updates. This section defines the data transmission requirements between the vehicle and the cloud. With a focus on AD/ADAS model training and service monitoring, it examines the implications for designing network bandwidth, onboard storage, and communication protocols required to support this lifecycle.

Note: The requirements discussed herein are derived from specific assumptions, including available vehicle sensor configurations and data-collection policies. These assumptions may vary significantly across automakers and software developers. Accordingly, the data volume and latency figures presented in the following sections should be regarded as illustrative examples rather than definitive benchmarks. They may be revised in future AECC publications as assumptions evolve, or new use cases emerge.

4.1 AD / ADAS Model Training

4.1.1 Service Scenario Description

AI models are a core element of modern Automated Driving (AD) and Advanced Driver Assistance Systems (ADAS). Inference tasks that enable road-situation awareness and motion planning are typically executed locally on vehicle Electronic Control Units (ECUs). In contrast, some data processing tasks can be offloaded to remote cloud/edge servers. The concept of Software-Defined Vehicles (SDVs) enables upgrading these in-vehicle AI models via over-the-air (OTA) software updates, allowing continuous improvement even after vehicles are deployed to the market.

Training AI models requires large volumes of data. In the context of AD and ADAS, sensor data generated by vehicles is the primary source for model development. This calls for a scalable data collection framework capable of transferring vehicle-generated data to cloud and edge computing environments, where data analytics and model training take place. Data analytics and AI model training are typically performed in batches, using datasets collected over a recent time window. As a result, data collection from vehicles can tolerate some delay, provided the required dataset is available before the next processing cycle begins. **The duration and frequency of these analytics cycles depend on the software development process, which may range from several days to several months**.

The software development team analyzes the collected dataset to retrain or finetune the AI inference models. This process typically involves GPU-accelerated data processing on public and/or on-premises cloud platforms. The result of model development is an updated software package that incorporates the refined AI inference models. This package must then be distributed to the relevant vehicles over the network within a defined timeframe — **ranging from several hours for critical updates to several weeks for non-safety-critical upgrades**.

4.1.2 Uplink Data Characteristics and Requirements

State-of-the-art ADAS and AD systems typically utilize a combination of cameras, radars, LiDAR, and/or other sensors. While the specific configuration and number of sensors vary across systems, for estimation purposes, we assume two representative sensor hardware setups.

AD with a high level of automation (e.g., SAE Automated Driving Level 3 or higher):

- 8 cameras
- radars

- 1 LiDAR
- Other sources, such as GNSS, IMU, and CAN bus

ADAS with a moderate level of automation (e.g., SAE Automated Driving Level 2):

- 6 cameras
- 5 radars
- Other sources, such as GNSS, IMU, and CAN bus

These assumptions do not reflect the current or future product specifications of specific automakers. They would rather define typical hardware configurations generally followed in the industry. To estimate data characteristics, we further make the following assumptions:

Data Recording Strategy: Instead of continuous recording, the system utilizes an intelligent trigger mechanism. High-fidelity sensor data is recorded only during critical "corner cases" (e.g., disengagements, near-misses, complex intersections)**.**

Recording Duration: Assuming **10 critical events per day**, with each event capturing **10 seconds of high-fidelity data (5-second pre-event buffer + 5-second post-event)**, the total high-quality recording time is 100 seconds per day per vehicle. Table 1 describes the example characteristics of the model training data for AD with a high level of automation. In contrast, Table 2 shows those for the ADAS with a moderate level of automation. Note that training data typically requires RAW or lossless formats to avoid compression artifacts that degrade model accuracy. In total, a vehicle generates **76.4 GB per day and 21.8 GB per day for AD and ADAS setups, respectively**, which need to be uploaded to the cloud/edge computing infrastructure for further analysis.

Table 1. Uplink Data Characteristics of AD Model Training (High Level of Automation)

Data Source	Quantity	Data Volume
Camera (Assumptions) 8.3 MP (4K resolution), 30 fps, 12 bpp, 1/4 compression rate **(lossless)** * MP: megapixels * fps: frames per second * bpp: bits per pixel	**8** Front: 2 Side: 4 Rear: 2	**74.7GB/day** 8 cameras * 8.3 MP * 1.5 bytes/pixel * 30 fps * 100 seconds per day * 1/4
Radar (Assumptions) Imaging radar: 30k pts/frame, 20Hz, 16 bytes/point, 1/4 compression rate (lossless) Standard radar: 2k pts/frame, 20Hz, 16 bytes/point, 1/2 compression rate (lossless)	**5** Front: 1 imaging radar Corner: 4 standard radars	**0.4 GB/day** Imaging: 1 sensor * 30k pts/frame * 20 Hz * 16 bytes/point * 100 sec/day * 1/4 = 0.24GB/day Standard: 4 sensors * 2k pts/frame * 20 Hz * 16 bytes/point * 100 sec/day * 1/2= 0.13 GB/day

Data Source	Quantity	Data Volume
LiDAR (Assumptions) 128 lines, 1.5M points/sec, 32 bytes/point including timestamp and reflectivity, 1/4 compression rate (lossless)	**1**	**1.2 GB/day** 1.5M points/sec * 32 bytes/point * 100 seconds / day * 1/4
Others (GNSS, IMU, CAN, etc.)	**Various**	**0.1 GB/day**

Table 2. Uplink Data Characteristics of ADAS Model Training (Moderate Level of Automation)

Data Source	Quantity	Data Volume
Camera (Assumptions) Front: 8.3 MP (4K resolution), 3p fps, 12 bpp, 1/4 compression rate (lossless) Surround/Rear: 2.1 MP (FHD resolution), 30 fps, 12 bpp, 1/4 compression rate (lossless)	**6** Front: 1 Side/Surround: 4 Rear: 1	**21.1 GB/day** Front: 1 camera * 8.3 MP * 1.5 bytes/pixel * 30 fps * 100 seconds per day * 1/4= 9.3 GB/day Others: 5 cameras * 2.1 MP * 1.5 bytes/pixel * 30 fps * 100 sec/day *1/4= 11.8GB/day
Radar (Assumptions) Front radar: 30k pts/frame, 20Hz Corner radar: 2k pts/frame, 20Hz 16 bytes/point 1/2 compression rate (lossless)	**5** Front: 1 standard long-range radar Corner: 4 standard medium-range radars	**0.6 GB/day** Imaging: 1 sensor * 30k pts/frame * 20 Hz * 16 bytes/point * 100 sec/day * 1/2= 0.5GB/day Standard: 4 sensors * 1k pts/frame * 20 Hz * 16 bytes/point * 100 sec/day *1/2= 0.06 GB/day
Others (GNSS, IMU, CAN, etc.)	**Various**	**0.1 Gbytes / day**

4.2 Service Monitoring

4.2.1 Service Scenario Description

After new software is rolled out, the developer should continuously monitor service operations to ensure they function as intended. In the event of anomalies—such as malfunctions or unexpected behavior—these issues must be detected promptly to enable rapid troubleshooting. Ongoing service monitoring is also essential for product improvement initiatives, such as A/B testing, where new features are deployed to a subset of customers, and their usage patterns and satisfaction levels are compared with those of customers using the current version.

Detecting critical anomalies is often time-sensitive, as delays in identifying and resolving service malfunctions can quickly erode customer satisfaction. To address this, vehicles should continuously analyze application and system logs and send alerts to the data-driven development platform whenever critical anomalies are detected. These time-sensitive notifications are best transmitted over cellular networks to ensure timely and reliable delivery.

In contrast, analyzing detailed service logs can help uncover hidden system issues and inefficiencies, which can be added to the development backlog and addressed in subsequent software update cycles. Data collection for this type of analysis is typically not time-critical, provided the logs are received before the next round of analytics begins. However, the volume of such log data is often much larger than that of time-sensitive anomaly alerts, as developers generally seek access to comprehensive records to better identify potential areas for improvement. This second type of data is well-suited to the Data-First Architecture, which enhances scalability by allowing vehicles to store and later offload large datasets opportunistically via Wi-Fi. Furthermore, to optimize bandwidth and storage, the system should support dynamic logging policies, enabling developers to remotely configure log verbosity and target specific vehicles for detailed tracing.

Typical examples of such service logs include the following:

- System & kernel logs (e.g., boot sequences, drivers, resource usage)
- Application logs (e.g., logic flow, errors, user interaction events)
- Telemetry data (e.g., sensor readings, CAN signals, geolocation – often the largest volume)
- Network & security logs (e.g., firewall events, connection latency)

4.2.2 Uplink Data Characteristics and Requirements

Vehicles run multiple applications and background services across various ECUs. The data volume fluctuates significantly based on driving conditions and debugging levels. To estimate the example data characteristics for storage sizing, we assume a high-fidelity data collection scenario:

- **Active applications & services**: A vehicle runs approximately 50 data-generating processes.
- **Data generation rate**: On average, the aggregate log generation across all systems is estimated at 500 MB per hour (uncompressed raw data).
- **Compression ratio**: Assuming efficient compression (e.g., Zstd), the data is reduced to approximately 100 MB per hour for storage and transmission.
- **Operation Time**: The vehicle operates for an average of 2 hours per day, but storage sizing should account for peak usage (e.g., long drives).

Based on these parameters, the daily log volume to be offloaded is approximately 200 MB per day under normal conditions. However, for high-definition data collection campaigns (e.g., ADAS sensor validation or comprehensive beta testing) and/or longer trips, volumes can easily reach 1 GB per day.

4.2.3 Downlink Data Characteristics and Requirements

In the monitoring ecosystem, downlink traffic primarily consists of software packages distributed via OTA. Once potential improvements are identified on the data-driven development platform, they are incorporated into upgraded software and deployed to vehicles. The size of these software update packages can vary significantly depending on factors such as application type, ECU hardware specifications, and optimization techniques (e.g., compression or differential updates). While edge ECUs require small updates, the shift to Zonal Architecture and centralized high-performance computer (HPC) significantly increases typical payload sizes. Modern SDVs utilize POSIX-based operating systems (e.g., Linux, QNX, Android), where a single system update can exceed 1 GB. Furthermore, high-definition map data and AI model updates constitute massive payloads. Therefore, distinguishing between “full image updates” for major version jumps and “delta updates” for routine patches is critical for bandwidth planning.

Table 3. Downlink data characteristics of the service monitoring use case

Software Type	Full Image Size	Delta Update	Frequency	Notes
Applications (Containers / Apps)	Tens of MB to Hundreds of MB	Several MB to tens of MB	High	Microservices or apps. Upgraded independently from the OS.
AI models (AD / ADAS)	Hundreds of MB to Several GB	Tens of MB to hundreds of MB	High	Frequent improvements to perception and planning models.
Infotainment Assets (Media, Graphics, Manuals, etc.)	Hundreds of MB to Several GB	Tens of MB to hundreds of MB	Medium	UI themes, voice packs, heavy media assets.
High Definition Maps	N/A	Several MB to tens of MB	Very high	Small tiles along planned/predicted routes are delivered *on demand*.
Standard Definition Maps	Several GB to tens of GB	Tens of MB to hundreds of MB	Low	Standard navigation maps.

5 High-Level Network Architecture

5.1 Layered Architecture

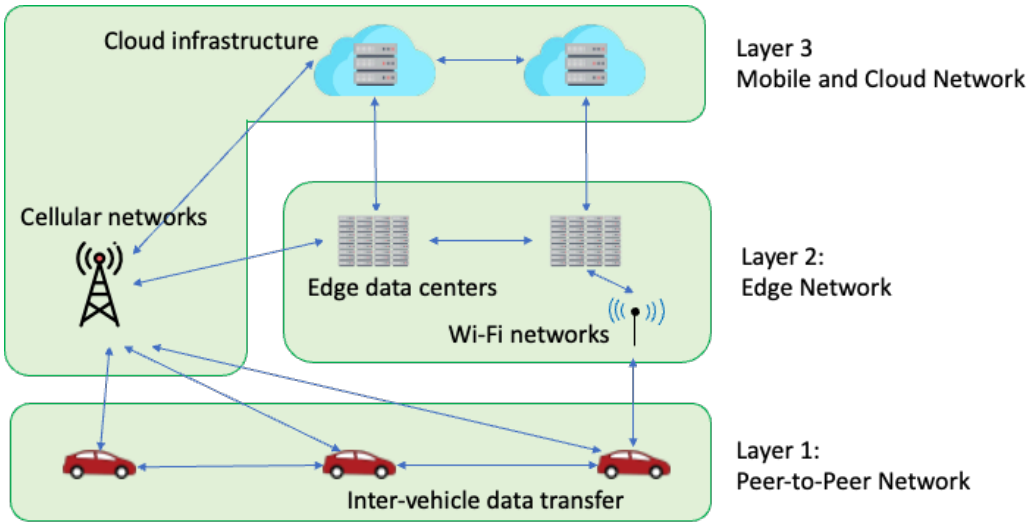

Figure 3: Layered architecture

The Data First Architecture comprises three network layers: peer-to-peer, edge, and mobile/cloud.

- **Peer-to-Peer Network Layer**: The first layer involves peer-to-peer communication links among vehicles, or more generally, among IoT devices. Distributing and aggregating data within local peer-to-peer networks reduces the number of active sessions, traffic volume, and server load on edge cloud infrastructure. A key component of this layer is a *data scheduler* that monitors vehicle status and connectivity with peers while efficiently scheduling data transmissions over peer-to-peer links (i.e., inter-vehicle data transfer).
- **Edge Network Layer**: The second layer incorporates short-range wireless communication infrastructure (e.g., Wi-Fi access points) and edge data centers. Its primary goal is to reduce the load on cellular radio access and core networks by opportunistically offloading traffic. Edge data centers should be designed with scale-in and scale-out capabilities to efficiently accommodate dynamic traffic demands.
- **Mobile and Cloud Network Layer**: The third layer consists of cellular networks and cloud computing infrastructure. This layer enables traditional automotive cloud computing platforms that integrate vehicles, mobile networks, and central cloud systems. It serves as the global layer for data distribution, aggregation, and lifecycle management.

The network architecture for software-defined vehicles should be designed to prevent performance bottlenecks and maximize scalability. Maintaining loose coupling among the three layers enables horizontal scaling, allowing the system to efficiently accommodate growing data communication and processing demands.

Energy efficiency is another critical aspect of a data-driven development platform. To support a sustainable ecosystem, systems should be designed to minimize power consumption for communications and data processing. In particular, reducing unnecessary data uploads, retraining jobs, and software or content distribution is essential for lowering energy usage and carbon dioxide emissions. This can be achieved through strategies such as selecting and sampling data at lower layers – such as within vehicles and peer-to-peer networks – before transferring it to higher layers, thereby reducing communication frequency and processing overhead. Additionally, upgraded content and AI models can be downloaded by only a small subset of vehicles and then disseminated within the peer-to-peer network layer. These optimizations improve the energy efficiency of both data communications and processing, contributing to a more sustainable development platform.

5.2 Data Transfer between Layers

Data transfers between layers must be carefully designed to maximize overall data management efficiency. For uplink communications from vehicles — particularly the collection of vehicle-generated data — the priority of each data item should be evaluated along two complementary dimensions:

- **Utility**: The operational and analytical value that the data provides, including its effective lifetime (TTL / freshness), its ability to improve analytics or model performance (e.g., informativeness and uniqueness), and its contribution to understanding rare or high-impact situations (e.g., penalty of data loss and relevance).
- **Constraints**: Unavoidable restrictions on how the data may be handled or propagated, such as privacy requirements, jurisdiction-specific rules, or policy-driven limitations on where the data may be transferred or retained.

Data with short lifetimes, high utility, and fewer constraints should be prioritized for transfer to upper layers. In contrast, lower-priority data can be retained within lower layers and discarded once its utility diminishes. To ensure reliable delivery of high-priority data, selective duplication within the peer-to-peer network layer may be employed, allowing a limited number of neighboring vehicles to forward the same data toward upper layers. The redundancy in data communication paths often facilitates reliable data delivery, while it does not always guarantee successful delivery to final destinations. The data reachability via inter-vehicle data transfer (e.g., data collection

among vehicles, followed by data offloading over Wi-Fi access points) is often lower than V2N communications, as policies on whether to upload other vehicles' data may vary across vehicles[1]

For **uplink communications** from a lower layer (e.g., a car) to an upper-layer resource (e.g., the cloud), the definition of the impact factor may vary across applications. In the context of AI model training, **example impact metrics** include:

- **Confidence**: When the transmitted data is the output of an inference model or algorithm (e.g., object detection), the confidence level of the estimation can serve as an indicator of the data's impact on the downstream model training task.
- **Informativeness**: Measures how much the sample reduces model uncertainty or improves decision boundaries when added to the training set.
- **Privacy**: It represents the mandatory constraints—legal, regional, and policy-based—that limit where data can flow and how it may be handled within the system.

Downlink communications from upper layers to vehicles follow similar prioritization criteria but with slightly different interpretations. In practice, downlink delivery must prioritize efficiency, favoring local caching and multi-hop dissemination (e.g., edge servers or inter-vehicle data transfer) to avoid unnecessary cloud-to-vehicle transfers and accelerate system-wide distribution.

- **Impact**: The degree to which the absence of the data would degrade the *quality of automotive services*.
- **Lifetime**: The duration over which the data remains useful for its intended purpose.
- **Relevance**: The proportion of vehicles for which the data is to be delivered.

Data with *greater impact*, *longer lifetime,* and *broader relevance* should be prioritized for transfer to lower layers. High-priority downlink data can be proactively disseminated within the edge network, for example, by caching it on edge servers as in content delivery networks, and peer-to-peer networks through inter-vehicle data transfer. These approaches accelerate distribution while adding redundancy to network paths, ensuring timely and reliable delivery.

Additionally, all data transfers must comply with regional privacy regulations such as GDPR, APPI, and CCPA. Interfaces between layers should include gateway mechanisms to block non-compliant transfers or apply necessary safeguards, such as anonymization or perturbation techniques.

5.3 Multipath Communications

5.3.1 Overview

Multipath communications are generally employed in applications with stringent requirements for latency and reliability. To address potential disruptions in active networks, existing solutions often utilize multiple cellular links from different mobile network operators and/or combine network interfaces of various types, such as Wi-Fi, non-terrestrial networks (NTN), and low-power wide-area (LPWA) technologies. By dynamically selecting network paths expected to deliver the best performance under given conditions, these systems can mitigate the impact of link degradation and maintain the desired quality of service. The delay-tolerant nature of network traffic in data-driven development platforms introduces additional flexibility in determining when to transmit data. Rather than sending data immediately over wide-area networks, vehicles can store it locally until they gain access to Wi-Fi

[1] Redundant forwarding in inter-vehicle data transfer is most effective when applied selectively only when needed.
Vehicles may exhibit different decision thresholds regarding data expiration and upload policies, leading to different data reachability characteristics between inter-vehicle data transfer and V2N communication paths. As such, redundant transmission improves reliability but does not guarantee reachability to upper layers.

access points or other local-area networks that offer data offloading. This approach reduces reliance on wide-area infrastructure and can improve overall transmission efficiency.

Figure 3 provides a high-level overview of multipath communications enhanced with scheduled transmission capabilities. The data generated by an application is passed to the *Multipath Communication Gateway* that manages transmissions from vehicles. Each data point is tagged with attributes such as delivery deadline and priority. The gateway comprises two key components: the *Data Scheduler* and the *Multipath Controller.*

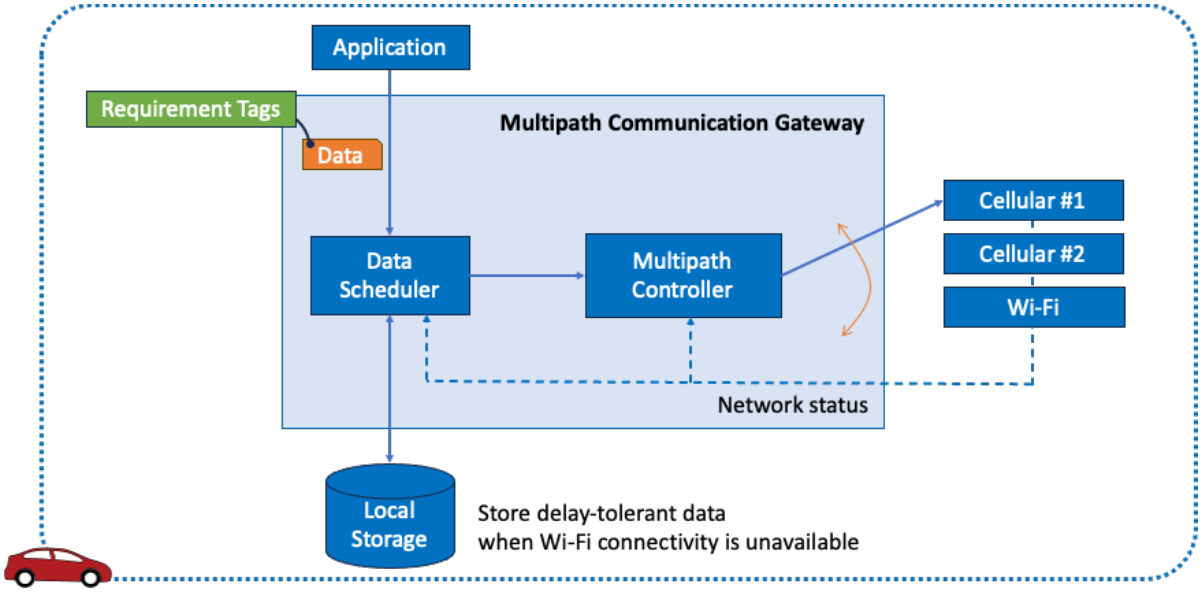

Figure 4: Multipath communications with scheduled transmission capabilities

5.3.2 Requirement Tags

When an application sends traffic to the Multipath Communication Gateway, it attaches requirement QoS (Quality of Service) tags to the data. These tags can be defined either per-packet or per-session, depending on the application type and the nature of the data traffic. The following is a list of classification items from the perspective of application data.

Examples of such tags include:

- **Deadline**: The time by which the data must be delivered to the intended recipient(s).
- **Priority**: The application-defined priority level of the data traffic. Since the definition of priority can vary between applications, it should be normalized into a unified format to enable consistent comparison across different sources.
- **QoS Class**: The minimum network quality level that must be guaranteed for application data. The required quality can be defined using a finite set of classes. Each class specifies performance requirements in terms of latency, reliability, bandwidth, and/or other relevant metrics.
- **Sensitivity Class**: Some application data traffic may contain sensitive information, such as personally identifiable information (PII). The transmission of such data may require special handling, for example, using specific communication paths (e.g., a cellular link provided by a designated mobile network operator) or restricting network routes to remain within certain geographical regions to comply with privacy regulations. An application data assigns one of the predefined Sensitivity Classes, each of which is linked to a corresponding set of network handling rules to be applied during transmission.

Sensitivity characterizes the inherent handling requirements of application data, defining the security, privacy, and regulatory constraints that govern how the data must be protected, authenticated, and shared across the system. Each sensitivity class can be associated with one or more of the following attributes:

- Required security level
- Permission for international data transfer
- Presence and/or type of privacy-sensitive information

5.3.3 Data Scheduler

The Data Scheduler evaluates each data item's delivery deadline. If sufficient time remains before the deadline, it may store the data locally and wait for Wi-Fi links to become available. It can also optimize cellular network transmissions by temporarily storing data during congestion and sending it in batches during off-peak hours to balance network load over time. The scheduler continuously monitors the remaining time; if the data cannot be transferred via Wi-Fi within the available window, it will be immediately sent over wide-area networks – such as cellular links – to ensure delivery delay requirements are met.

5.3.4 Multipath Controller

Once the scheduler determines that data should be transmitted immediately, the Multipath Controller in the Multipath Communication Gateway selects the network interface(s) to use for transmission. The controller may choose either a single network interface expected to deliver optimal performance or multiple interfaces to enable parallel transmissions for increased throughput or reliability. In making this selection, the Multipath Controller evaluates the requirement tags attached to each data item to ensure compliance with the specified Quality of Service (QoS) parameters and any applicable regulatory constraints.

5.3.5 Optimization of Individual Networks and Data Connections

In addition to data scheduling and communication path selection, it is essential to fine-tune individual networks to better support vehicle mobility. For Wi-Fi communications, reducing link establishment latency directly affects the amount of data that can be transferred over transient connections between vehicles and Wi-Fi access points. When signal quality indicators—such as received signal strength—begin to degrade, or when a station detects severe channel congestion on active links, proactively disconnecting from the current access point and searching for an alternative may enhance overall data offloading performance.

It is also important to control the number of active data connections per vehicle. Rather than allowing individual applications to establish network sessions independently, a coordinating function can manage and aggregate the sessions originating from the vehicle. Additionally, the frequency of Transport Layer Security (TLS) negotiations should be minimized to reduce the overhead associated with certificate exchanges.

6 Design Considerations

6.1 Radio Access to Wi-Fi Access Points

6.1.1 Initial Link Setup Latency

Communication links between vehicles and roadside Wi-Fi access points are typically transient, especially when vehicles are in motion. For example, if a vehicle is traveling at 60 km/h and the Wi-Fi access point along the road has a communication radius of 100 m, the vehicle has only about 12 seconds to exchange data. In practice, the effective range is often much shorter due to radio interference and signal attenuation from buildings and other obstacles. Wi-Fi access points near parking areas are generally more promising for data transfer, as vehicles move slowly during parking maneuvers. However, vehicle ECUs are typically powered down shortly after the ignition is turned off. This means that data offloading must be completed within a short window—starting when the vehicle enters the access point's communication range and ending before the ECU powers down. To maximize the amount of data offloaded via Wi-Fi, vehicles should minimize the latency of initial link setup.

The latency in establishing an initial Wi-Fi link is primarily influenced by the following factors:

- **Access Point (AP) discovery** – The Wi-Fi station scans each radio channel to identify nearby APs.
- **Authentication and association** – The station and AP exchange a series of control messages to establish a secure communication link.
- **IP address assignment** – A DHCP server behind the AP allocates a dynamic IP address to the newly connected station.

In a naïve implementation, the entire process—from AP discovery through IP address assignment—can take from several seconds to more than 10 seconds, consuming a substantial portion of the available Wi-Fi link lifetime.

AP discovery latency can be reduced by leveraging prior knowledge of the operating channels used by known Wi-Fi access points (APs). A service provider can maintain a database of AP operational status, collected from Wi-Fi controllers managed by network operators. Vehicles can query this database and scan only the radio channels associated with APs located near their current position. This targeted scanning approach significantly reduces AP discovery time compared to the naïve method of scanning all radio channels supported by the Wi-Fi firmware and permitted by local regulations.

The latency of authentication, association, and IP address assignment can be reduced by utilizing the **Fast Initial Link Setup (FILS)** feature, as defined in the IEEE 802.11ai amendment to the IEEE 802.11 specification [3]. FILS significantly reduces the number of control messages required for Wi-Fi authentication while simultaneously performing IP address assignment during the authentication and association process. This streamlined procedure can greatly shorten overall link setup time.

Even short connection durations can be sufficient for meaningful data exchange if the link is established promptly. The following examples illustrate that minimizing Wi-Fi connection and authentication latency directly expands the range of practical automotive use cases, even for short-term or temporary connectivity opportunities.

Representative examples include:

- **Low-speed driving (≈30 km/h)**: When a vehicle passes through a coverage radius of approximately 50 m, the available connection time is around 12 seconds. At an effective throughput of 15 Mbps, this allows uploading roughly 22.5 MB of data, which can be sufficient for transferring diagnostic logs or monitoring data.

- **Deceleration to stop (≈18 km/h)**: Under gradual deceleration, the same coverage distance yields an interaction window of approximately 20 seconds, enabling larger batches of telemetry or partial content updates to be exchanged.
- **Stopped or ultra-low-speed driving (≈6 km/h)**: When the vehicle is stopped or moving very slowly, the available connection time can extend to approximately 60 seconds. At an effective throughput of 20 Mbps, up to 150 MB of data can be downloaded, making this scenario suitable for map segment updates or other data-intensive downloads.

6.1.2 Coverage Extension by Mesh Networking

The effective communication range of a Wi-Fi access point is typically less than 100 m, which is often insufficient to cover large parking lots or parking spaces with numerous obstacles. When multiple vehicles are located within the same parking area, they can form a Wi-Fi mesh network to extend effective communication coverage. Some Wi-Fi product manufacturers have implemented proprietary extensions to the IEEE 802.11s mesh networking specification. At the same time, the Wi-Fi Alliance has defined the Wi-Fi EasyMesh™ specification to improve interoperability among vendors. Leveraging inter-vehicle mesh networking can provide data offloading opportunities to many vehicles that would otherwise have no direct connectivity to Wi-Fi access points.

6.2 Wi-Fi as an Enabler of Inter-Vehicle Data Transfer

6.2.1 Wi-Fi-based Peer-to-Peer Communications

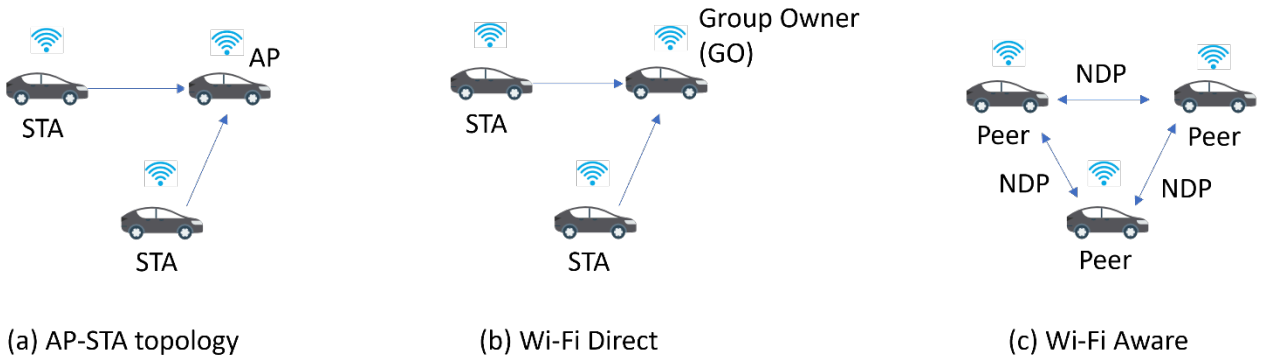

Figure 5: Example embodiments of Wi-Fi-based inter-vehicle data transfer

Vehicles in motion can use direct Wi-Fi communication links between vehicles for inter-vehicle data transfer. There are three possible embodiments of Wi-Fi peer-to-peer communications, each having advantages and limitations for inter-vehicle communications, as summarized in Table 4.

Table 4. Comparison of Wi-Fi peer-to-peer communication technologies

Technology	Topology / Operation	Security	Link Setup Latency	Topology Management
AP-STA Topology	Star topology where some vehicles act as Wi-Fi Access Points (APs) and others connect as stations (STAs).	**Advantage:** Supports strong Wi-Fi security, including **WPA3 Enterprise**	**Limitation:** AP discovery can be slow because stations typically scan all radio channels	**Limitations:** • AP vehicles cannot communicate directly with other AP vehicles • Requires an external mechanism to decide which vehicles become APs
Wi-Fi Direct	Peer-to-peer mode with a **Group Owner (GO)** acting as a temporary AP. Other devices join the GO's group. GO is selected either pre-configured or via decentralized negotiation.	**Limitations:** • Supports only **WPA2 Personal** • Uses **Wi-Fi Protected Setup (WPS),** requiring user interaction (e.g., PIN or button press), unsuitable for inter-vehicle data transfer	**Advantage:** Peer discovery occurs on predefined social channels, reducing channel scanning and discovery latency	**Advantage:** Decentralized GO negotiation automatically selects a single coordinating node
Wi-Fi Aware	Fully decentralized peer-to-peer operation. Devices discover each other via service advertisements and establish direct **NAN Data Paths (NDPs)** without APs or GOs.	**Limitations:** • Supports only **Pre-Shared Key (PSK)** authentication • No enterprise-grade security support	**Advantage:** Peer discovery occurs on a single channel, minimizing discovery time	**Advantage:** No centralized coordinator required; pure peer-to-peer topology

Although several standards exist for Wi-Fi-based peer-to-peer communications, each has its own limitations when applied to inter-vehicle communication scenarios. Automotive variants of IEEE 802.11, such as IEEE 802.11p and IEEE 802.11bd, have also been established; however, these are optimized for safety-related applications comprising the Intelligent Transportation Systems (ITS) bands. They are not well-suited for reliably transferring large volumes of private data, as envisioned in the service scenarios discussed in Section 4. This situation underscores the need for an industry-wide discussion of potential enhancements to existing wireless LAN standards to better support emerging automotive use cases.

6.2.2 Authentication for Inter-Vehicle Data Transfer

Security is always a top priority in any communication system, and Wi-Fi-based inter-vehicle data transfer is no exception. However, unlike infrastructure-mode Wi-Fi, authentication mechanisms for peer-to-peer communication have not yet been well-established or standardized. From a security perspective, mutual TLS-based

authentication would be a preferable approach, as it enables not only the AP (or equivalent coordinating node) to verify connecting stations, but also allows the stations to authenticate the AP they are attempting to connect to. This two-way authentication is particularly important in peer-to-peer environments, where trust must be established in both directions to prevent malicious impersonation or unauthorized data access.

A possible approach to inter-vehicle Wi-Fi authentication is to deploy a remote authentication server on a cloud or edge server that vehicle-based APs can access via V2N (e.g., cellular) communications. This architecture extends the typical authentication mechanism for enterprise Wi-Fi networks, which rely on wired connectivity for authentication server access, to inter-vehicle communication scenarios. However, this approach presents the following drawbacks:

- **Authentication latency**: Wi-Fi authentication involves multiple message exchanges between stations and authentication servers via APs, resulting in non-negligible latency. This latency becomes problematic when communicating with remote authentication servers, as vehicle-based APs typically maintain cellular or satellite connections for backhaul. In typical enterprise Wi-Fi deployments with wired connectivity, this delay is often manageable. However, the extended authentication latency in inter-vehicle communication scenarios reduces the number of communication opportunities, as vehicles may move out of range before establishing a secure connection.
- **Overhead**: The centralized authentication architecture requires interactions with the remote authentication server every time a vehicle attempts to connect to another vehicle in the vicinity. This results in a substantial volume of authentication requests from vehicles to the authentication server, increasing V2N communication overhead and server load.
- **Availability**: Vehicles cannot establish an inter-vehicle Wi-Fi communication link outside the coverage area of V2N communication networks (e.g., cellular base stations). This represents a significant limitation, as a major benefit of inter-vehicle data transfer is reduced reliance on network infrastructure availability.

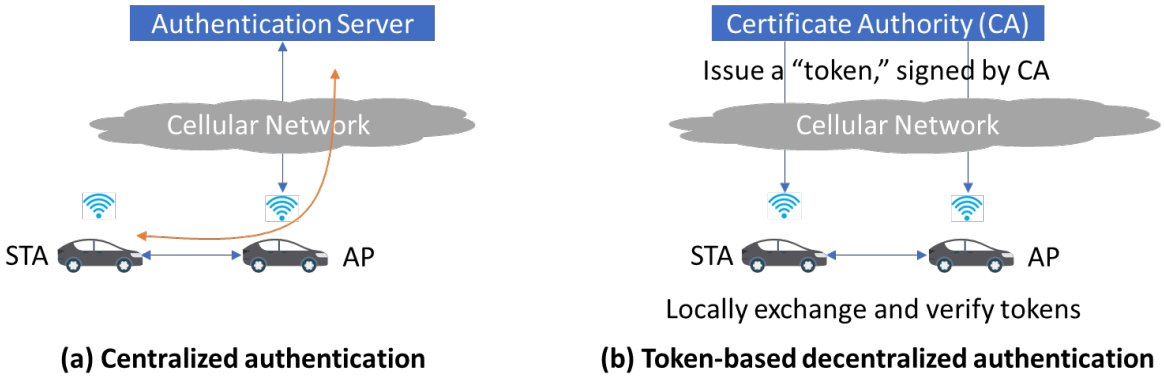

Figure 6: Inter-vehicle Wi-Fi authentication

Addressing these challenges requires extensions to the existing authentication mechanisms for enterprise Wi-Fi networks. A possible solution is **token-based authentication**, in which a trusted certificate authority issues a token to individual vehicles that wish to join inter-vehicle Wi-Fi networks. The token includes the vehicle's identity, along with a set of conditions for establishing an inter-vehicle communication link, such as the group of vehicles it is allowed to connect to, a communication rate limit, and security requirements. The certificate authority digitally signs the token, which becomes valid for a limited period (e.g., 30 minutes). When a vehicle attempts to establish a connection with another vehicle, the vehicles exchange their tokens over the inter-vehicle Wi-Fi link and verify the token of the connecting peer. The connection is established only if (1) the token of the connecting peer is valid and (2) the conditions declared in the token are met.

A key advantage of the token-based approach is that the authentication process can be completed locally between a pair of connecting vehicles without communicating with a remote authentication server. This reduces the authentication delay associated with V2N communication latency while mitigating the dependence on cellular network coverage. Each vehicle should access the certificate authority only when the current token expires to get its token renewed. The certificate authority can be hosted within a cellular core network to leverage SIM-based authentication managed by MNOs, or on edge-based or Internet-based servers operated by third-party providers.

Some service scenarios may require inter-vehicle Wi-Fi communication between vehicles from different OEMs and/or fleet operators. For instance, public transport providers could offer a “data ferry” service in which fleet vehicles (e.g., taxis, buses) act as mobile Wi-Fi access points. These vehicles would collect data from other vehicles and IoT sensors along their routes, then upload it via access point infrastructure near parking areas. Standardizing token formats and federating certificate authorities across different stakeholders (such as public transport operators and multiple automotive OEMs) would enable seamless cross-operator collaboration in automotive data collection.

7 Conclusion and Next Steps

The transition to SDVs has fundamentally altered the automotive landscape, establishing a continuous "DevOps loop" where vehicle data drives the evolution of AI models and software features. However, the exponential growth of automotive big data—ranging from high-fidelity sensor data for AD/ADAS training to comprehensive service logs—has outpaced the capacity and cost-efficiency of traditional cellular-centric IoT architectures.

In this white paper, the AECC has proposed the **Data-First Architecture**, a scalable framework that places data at the core of the ecosystem. By implementing a hierarchical computing model that integrates Peer-to-Peer, Edge, and Mobile/Cloud networks, this architecture transcends the limitations of single-path connectivity. We have demonstrated that by leveraging **Multipath Communications** and intelligent **Data Scheduling**, vehicles can effectively manage delay-tolerant traffic. This allows offloading massive datasets via Wi-Fi and inter-vehicle communication links while reserving cellular bandwidth for time-sensitive, critical operations. Furthermore, this approach promotes sustainability by optimizing energy consumption through reduced transmission frequency and localized data processing.

Next Steps and Industry Challenges

To fully realize the vision of the Data-First Architecture, several technical and organizational challenges must be addressed through industry-wide collaboration:

- **Standardization of Inter-Vehicle Wi-Fi Authentication:** As discussed in Section 6, scalable inter-vehicle communication requires a robust security framework that does not rely on constant cellular connectivity. The industry shall work to standardize decentralized, token-based authentication mechanisms and federate certificate authorities. This will enable secure "data ferry" services and seamless cooperation across different OEMs and fleet operators.
- **Enhancement of Wireless Standards:** While current Wi-Fi standards offer peer-to-peer capabilities, they present limitations regarding connection setup latency and security in automotive scenarios. Continued dialogue is required to refine these standards or develop enhancements to better support the rapid, high-volume data exchanges required by moving vehicles.
- **Cross-Domain Ecosystem Development:** The shift from exclusive, single-operator models to collaborative resource sharing requires new business frameworks. Stakeholders must explore deployment patterns that incentivize the rollout of edge infrastructure and shared access points, ensuring that the benefits of data-driven development are accessible across the connected vehicle ecosystem.

The AECC remains committed to refining these architectural principles and driving the technical standards necessary to support the next generation of connected services.

8 References

[1] Automotive Edge Computing Consortium, “General Principles and Vision,” Version 4.0.6, Feb. 2026. [Online]. Available: https://aecc.org/wp-content/uploads/2026/02/AECC_General_Principle_and_Vision_v4.0.6_Feb-26_2026-1.pdf

[2] IDC, "Future Trends in Connected Vehicle Data Monetization," IDC Market Analysis, 2023.

[3] IEEE, "IEEE Standard for Information Technology - Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," in IEEE Std 802.11-2020 (Revision of IEEE Std 802.11-2016), pp.1-4379, 2021.