AECC
AUTOMOTIVE EDGE COMPUTING CONSORTIUM

SBD

WHITE PAPER

Coordinating the Connected Car and Network Edge (... to Avoid Breaking the Internet)

Today's internet is not ready for tomorrow's connected cars

Future automotive workloads are pushing performance boundaries at the Network Edge, potentially to breaking point. The connected vehicle ecosystem is rapidly transforming from delivering a basic telematics service to a combination of data-rich consumer services and data-intensive transportation-enabler services. Thirty-one million vehicles with embedded forward-facing cameras will be sold in Western markets in 2025, generating up to 10 Exabytes of data per day to be processed by off-board resources. According to the Automotive Edge Computing Consortium (AECC), network sustainability will become an issue if the majority of that data is all being uploaded in peak times (Mobile World Live, 2022). Without significant change to both vehicle data management and telecommunications network infrastructure plans, and coordination between the two, the automotive industry looks set to break the internet by 2030.

31M

31 MILLION VEHICLES WITH FORWARD-FACING CAMERAS WILL BE SOLD IN WESTERN MARKETS IN 2025 (SBD AUTOMOTIVE, 2022)

Today, typical IoT device communication patterns reflect a download-centric internet. They send a low-payload request from an app to a server, which sends a load of data back. Whether streaming movies, downloading album art, or gaming, the upload burden is significantly lower than the download. The internet is designed to support this pattern, with download speeds typically operating at approximately three to ten times the upload speed (Speedtest, 2022). Speedtest's 2022 Global Index shows that in the USA in 2022, mean broadband download speeds were nearly three times the upload speed (1225Mbps vs. 86Mbps). Mobile speeds show an even greater gap (128Mbps v 14Mbps). Connected, and increasingly autonomous, vehicles, follow the opposite pattern, which is a problem for mobile network providers.

128MBPS

128MBPS MEAN FIXED MOBILE DOWNLOAD SPEED IN THE U.S. DWARFS THE 14MBPS MEAN DOWNLOAD SPEED

The data coming from cars to support autonomous driving models in the cloud include extensive amounts of camera data to perform tasks such as HD mapping, road surface monitoring, or vehicle security monitoring. Without better coordination of data flow among in-vehicle, near-vehicle (network edge), and cloud compute resources, the internet usage patterns a vehicle manufacturer anticipates for the future will require so much network infrastructure investment that it becomes infeasibly expensive. The advances of 5G to help individual base stations operate faster and reduce backhaul requirements compared to 4G will be dwarfed by data being sent to places it perhaps does not need to go.

Industry PoCs Support Advanced Development

Partners from across the connected vehicle industry are testing these data-heavy use cases in real driving conditions to understand how the enabling infrastructure needs to evolve to cope with the data volumes that make advanced features possible. Some early network edge proofs-of-concept (PoCs) focused on very local hazard identification and communication use cases, primarily around intersections (HERE Technologies, 2020) (Vodafone, 2019).The primary goal was to validate the expectation that device-to-device communication via a local Multi-access Edge Computing installation (MEC) could quickly and efficiently warn drivers or vehicles of unseen hazards such as pedestrians, cyclists, or emergency vehicles. MECs tend to be deployed at local level to base stations or local servers supporting a small number of base stations, as opposed to centralized cloud centers, making them ideal resources for services that are only locally relevant. The local nature of these use cases meant that the question of scale – whether data volumes or connected units – did not form a part of the early automotive-edge studies.

More recently, studies supported by the AECC have begun looking at how vehicles and the network might have to adapt to handle the scale of image-laden workloads, including storage, that would swamp the network infrastructure as it stands today. Table 1 shows how the cross-industry focus is changing to look towards future infrastructure concern points and opportunities.

Network Impact Considerations

	Scalability (# Connections)	Latency	Bandwidth	Localized Workload	Distributed Workload	Moveable Workload	Workload Prioritization
AECC (Oracle, AWS, Microsoft, Google, Wasabi)	✓	✓			✓		
AECC (Vodafone + KDDI)	✓					✓	✓
AECC (Vodafone, Ericsson, Toyota)	✓						✓
AECC (KDDI, DENSO, Toyota)	✓			✓	✓		✓
AECC (KDDI + Nexar + Oracle)		✓		✓			
Vodafone + Continental		✓		✓			
Verizon + Nissan		✓		✓			
Verizon + LG + AWS + Renovo		✓					
HERE + Verizon		✓	✓	✓			

Table 1: Primary technical impacts considered in recent PoCs

Managing the Complexity

Future data-intensive use cases are driving the transition from short-range, low-bandwidth data studies, which mainly focus on local C-V2X safety features like detecting hazards around a corner, to studies that consider how to use network edge facilities as a key part of a vehicle-to-cloud infrastructure. Driverless vehicle teleoperation when the vehicle enters an area or situation where it is not able to safely drive itself, for example, will involve significant data upload which requires an ultra-low latency in terms of the teleoperator's response. In the near term, such cases will be relatively few as driverless vehicle fleets remain small and teleoperation may be focused on the freight industry. But the fleets will grow, increasing the importance of balancing balance critical and non-critical network traffic.

A less futuristic use case involves coordinating detailed live map information such as relative positions of moving objects like people and vehicles with static maps to create a live high-definition (HD) map. HERE and Verizon's 2020 study included image processing on the edge to create a Visual Positioning Service that provides 3D positioning to complement satellite positioning technologies – which can be very helpful in multi-story parking facilities or multi-layer road systems (HERE Technologies, 2020). The volume of data to perform such calculations at scale could very easily overrun local fixed network infrastructure if the processing is in a central location. On the network edge, though, the scale becomes manageable.

To understand how even basic image processing in the cloud can generate worrying levels of network traffic, it helps to consider this use case across several large metropolitan areas. Tracking the location and sensor data at peak times to support such a function will require active data connections for millions of devices. For example, the Nagoya area of Japan was home to approximately 1.1 million privately-owned passenger vehicles in 2021, or 1 for every two people (Statista, 2022). A reasonable estimate for rush hour in Nagoya would see 50% of people on the move, with 20% of them in vehicles. One hundred thousand vehicles moving along with 550 thousand phones, all potentially contributing to a live map in the cloud which requires the map to be shared among the 100,000 vehicles, would put a considerable load on the network. Replicating that across a region begins to pressurize network infrastructure.

Building a future network that can capably cope with dozens of such use cases in the cloud, from a range of service providers and vehicle manufacturers requires managing upload volumes at a level contrary to the way the infrastructure is designed. Therefore, the automotive and network industries' joint trials are critical to specifying how the network edge layer can increase off-vehicle processor volumes without breaking the internet.

Long-term automotive workloads at the network edge face significant challenges

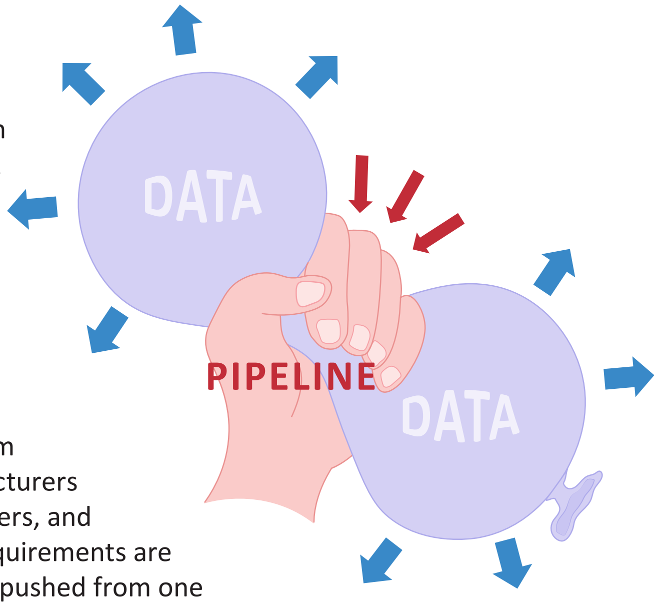

Recent proofs-of-concept trials performed by AECC build on the body of cross-industry work showing **what** the network edge can do. These PoCs focus on **how** vehicle data traffic between the vehicle and the network edge can be managed to preserve low-latency for critical services and local services while supporting end-to-end processing for important time-insensitive services. These PoCs from AECC have developed beyond early cross-industry trials of 3GPP / 5G protocols. The latest trials are validating requirements and specifications that have been created with support from stakeholders in all the relevant industries – vehicle manufacturers and suppliers, network operators and infrastructure providers, and cloud infrastructure and service providers. Data volume requirements are like a balloon - the volume does not change when the air is pushed from one side to another. Cross-industry approaches that aim to press the data balloon from all sides until it is the right shape to efficiently flow through the network are critical to delivering the combination of exciting features and ultra-low latency services that the future connected vehicle offers.

A study performed by KDDI, DENSO, and Toyota focused on what the AECC call "Opportunistic Data Transfer (ODT) (Mobile World Live, 2022). ODT involves prioritizing service data so that time-insensitive data destined for the cloud can be stored locally and transferred when the network is not busy. The concept builds on the Background Data Transfer protocol standardized by 3GPP and requires on-vehicle storage capacity to cache non-critical sensor data for upload outside peak network hours. Some vehicle manufacturers may object to the idea of adding 50-100MB of RAM to their vehicles, but if the alternative is to pay network fees sufficient to fund the investment in base stations, MECs, and backhaul to support these high-bandwidth features, the cost of chips may be easier to bear. From the network provider's view, a 90-9-1 data management concept offers a guideline to carmakers about how to maximize the efficiency of their connected services and data collection.

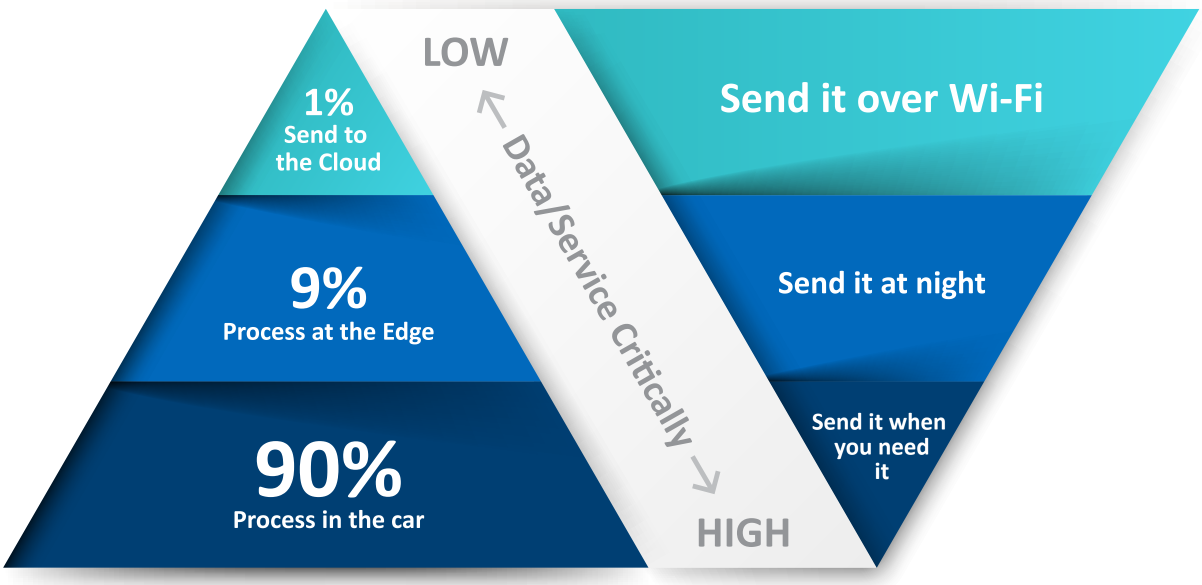

Figure 1: Upload-based service management on a 90-9-1 basis is analogous to the

approach many carmakers take to over-the-air software updates

Aiming for data processing that is "90% in the car, 9% on the edge, and 1% sent to the cloud" strikes the right balance to deliver mission-critical, low-latency services, says AECC board member Christer Boberg, Director of Ericsson's Group Strategy. Many carmakers already implement some level of on-board data aggregation or off-peak software downloads to reduce their network transmission costs. Applying this principal to sensor-data uploads to achieve a 90-9-1 breakdown may require changes to the on-vehicle storage and processing capacity long-term plans, which is why network operators, vehicle manufacturers, and automotive suppliers are working together to validate the specifications for both prioritizing the data and communicating end-to-end to move the data at the most appropriate time.

A PoC from KDDI, Nexar, and Oracle shows a different approach to prioritizing data flows by localizing the heavy image processing to precisely the edge node(s) that will deliver the service with the cloud servers providing the overall service – distributing the workload based on the level of local, compared to regional, relevance. This study involves using vehicle dash camera data to identify and pinpoint on a map the available parking spaces using a nearby MEC, with a cloud center supplying the availability data (KDDI Corporation, Nexar Inc., and Oracle Japan Corporation, 2022).

As Figure 2 shows, such a use case allows the MEC to filter out the high-volume data, leaving the network and cloud infrastructure to handle the processing outputs instead of the raw input.

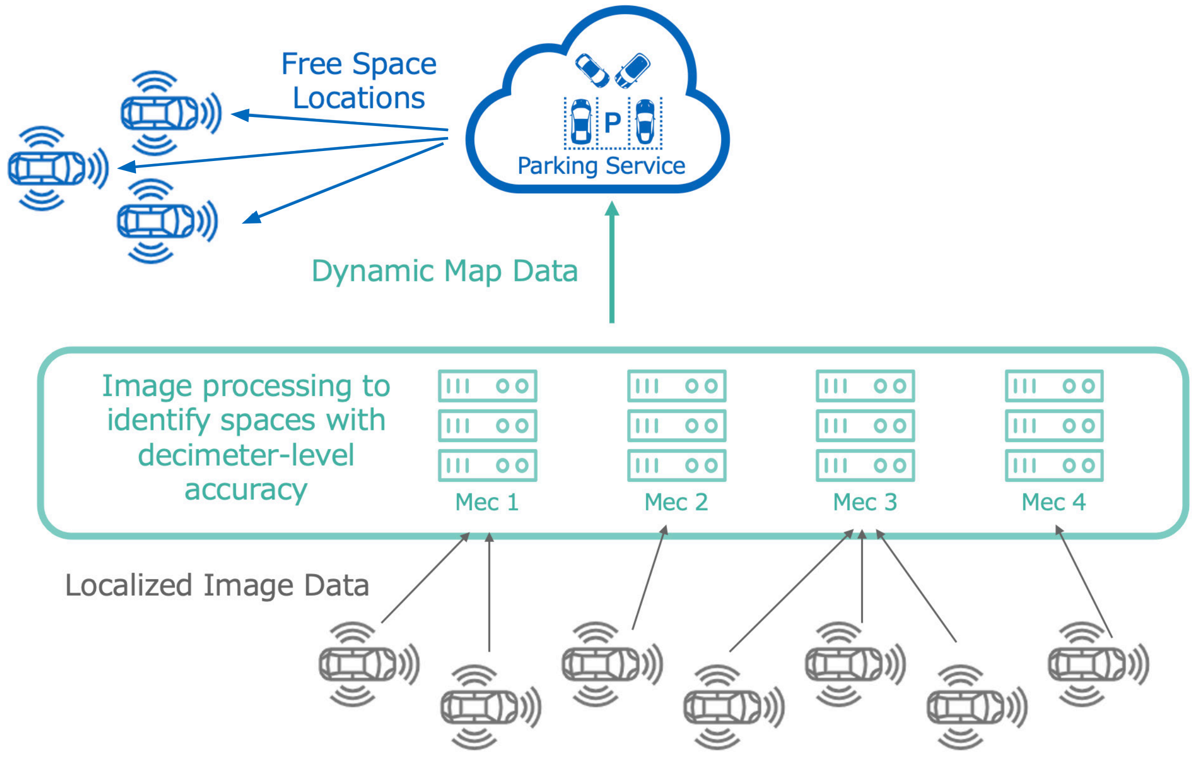

Figure 2: Local MECs process image and location near vehicles, reducing the data volume to the cloud service infrastructure

In this case, the data from each vehicle driving on a street may not be large, and there may not be a vast number of vehicles sending data at any one time. However, performing the image data processing at a regional level would result in large amounts of data transmission within the network that are only relevant to a very specific location. For example, people in New York City rarely need to have detailed parking information for Philadelphia. Within New York City, the parking maps for different parts of Manhattan are unrelated to parking maps for Queens. Creating a live map of available spaces on a small number of streets with a single MEC keeps the bulk of the data processing as local to the parking spaces as the infrastructure allows, reducing network traffic that would otherwise go to a regional or centralized cloud.

These PoCs are aimed at validating AECC's specifications so they can be rolled out in the real world. Gradually, more studies will stress-test the work of the consortium and its members. But beyond trying to ensure the data balloon can pass through the network pipe.

What does the industry need next: looking beyond data volume concerns?

Beyond the basic concerns about preventing future connected vehicles from breaking the internet, connected vehicle stakeholders have other challenges to deal with in providing the future HD mapping, image-processing, and ultra-low latency features that industry PoCs have been testing so far.

A key upcoming question for the connected vehicle industry is, "how can these exciting features be delivered consistently across a vehicle fleet or within a market?" If so much of the functionality relies on local processing of local data, car manufacturers need to understand how the same feature performs in Des Moines and San Francisco, or Frankfurt and Aachen. Network infrastructure is rarely deployed uniformly in a market, so managing differences in geographical capability or performance will present hurdles that need to be addressed.

Equally, one of the main reasons carmakers want to process data in the cloud is to pay once for the infrastructure in a centralized location instead of millions of times in vehicles around the globe. The 90-9-1 concept works well for network providers but may add a cost burden to some vehicle segments that is unsustainable. Working together to find flexible ways of re-shaping the data balloon can help vehicle manufacturers to have a range of data-division approaches to suit their full vehicle fleets.

Today, even the most advanced markets do not have the uniform network infrastructure to support the automotive industry's ambitions. Cooperation among the different industries that enable vehicle connectivity is key to addressing the data pipeline and processing investment constraints as well as other questions that will arise as specifications that support data volume balancing mature and the automotive world begins to look at deploying these data-intensive services at scale.

References

HERE Technologies. (2020, January 06). *Verizon and HERE join forces to create revolutionary safety and navigational systems.* Retrieved from HERE Technologies: https://www.here.com/company/press-releases/en/2020-06-01-0

KDDI Corporation, Nexar Inc., and Oracle Japan Corporation. (2022, August). *Enabling a Geolocation Parking Service with AECC Distributed Edge Architecture.* Retrieved from AECC: Proof of Concepts: https://aecc.org/wp-content/uploads/2022/08/AECC_-_Geolocation_PoC_Designed_Final_8.9.pdf

Mobile World Live. (2022, 06). *Connecting to Opportunity: AECC shows PoCs for next¬gen vehicles.* Retrieved from AECC: https://aecc.org/wp-content/uploads/2022/06/MWL_-_AECC_whitepaper_-_Connecting_to_Opportunity_E28093_AECC_shows_PoCs_for_nextgen_vehicles.pdf

SBD Automotive. (2021). *Connected Car Forecast 2021.*

SBD Automotive. (2022). *EU ADAS and Autonomy Forecast 2022.*

SBD Automotive. (2022). *USA ADAS and Autonomy Forecast 2022.*

Speedtest. (2022, June). *United States Median Speeds June 2022.* Retrieved from Speedtest Global Index: https://www.speedtest.net/global-index/united-states

Statista. (2022, February). *Leading cities with the largest number of private passenger cars owned in Japan as of March 2021.* Retrieved from Statista: https://www.statista.com/statistics/1292161/japan-cities-with-most-passenger-cars/

Vodafone. (2019, February 22). *Use of Multi-access Edge Computing to power artificial intelligence for automotive.* Retrieved from Vodafone: https://www.vodafone.com/news/technology/multi-access-edge-computing-to-power-artificial-intelligence-for-automotive

About the AECC

The Automotive Edge Computing Consortium (AECC) is an association of cross-industry, global leaders working to explore the rapidly evolving and important data and communications needs involved in instrumenting billions of vehicles worldwide. The AECC's goal is to find more efficient ways to support the high-volume data and intelligent services needed for distributed computing and network architecture and infrastructure. The AECC's members are key players in the automotive, high-speed mobile network, edge computing, wireless technology, distributed computing, and artificial intelligence markets.

If you would like to learn more, **visit AECC.org** or **contact the AECC here**.

About SBD Automotive

SBD Automotive is a global consultancy firm specializing in automotive technologies. For over 20 years, our independent research, insight, and consultancy has been helping vehicle manufacturers and their partners to create smarter, more secure, better connected, and increasingly autonomous cars. With a reputation for robust data and expert advice, as well as an ability to attract and retain the industry's most talented specialists, SBD Automotive operates a network of local offices in key automotive hubs, including the UK, Japan, USA, Germany, China, and India.

If you would like to learn more, please **email SBD Automotive**.