AvidThink®

Sponsored by aws

Mobile Edge in Focus - Automotive

Exploring Edge Computing Benefits for the Automotive Industry

RESEARCH BRIEF

Table of Contents

Research Briefs are independent content created by analysts working for AvidThink LLC. These reports are made possible through the sponsorship of our commercial supporters. Sponsors do not have any editorial control over the report content, and the views represented herein are solely those of AvidThink LLC. For more information about report sponsorships, please reach out to us at **research@avidthink.com**.

About AvidThink®

AvidThink is a research and analysis firm focused on providing cutting edge insights into the latest in infrastructure technologies. Formerly SDxCentral's research group, AvidThink launched as an independent company in October 2018. Over the last five years, over 110,000 copies of AvidThink's research reports (under the SDxCentral brand) have been downloaded by 40,000 technology buyers and industry thought leaders. AvidThink's expertise covers Edge and IoT, SD-WAN, cloud and containers, SDN, NFV, hyper-convergence and infrastructure applications for AI/ML and security. Visit AvidThink at **www.avidthink.com**.

RESEARCH BRIEF

Mobile Edge in Focus – Automotive

Exploring Edge Computing Benefits for the Automotive Industry

Overview

In this series of research briefs, AvidThink will examine mobile edge computing opportunities across select enterprise verticals. As we laid out in our 2020 Edge and Beyond report, edge computing spans multiple locations, with different tradeoffs at each. Likewise, for each vertical, the benefits provided by the edge will differ. This brief will focus on how the automotive industry can benefit from mobile edge computing.

For CIOs, technologists, developers, and business executives in the automotive industry, we will provide an overview of mobile edge computing, explain its general benefits, and discuss how automotive trends align with what the edge can offer. We will discuss promising use cases for edge technology in the automotive sector and then wrap up with a concrete set of steps to get started.

What is Mobile Edge Computing?

Today's public and private enterprise clouds are hosted within large data centers. While these data centers benefit from economies of scale and provide application developers with a rich set of supporting services, there are use cases where keeping computing resources close to users and enterprises can be beneficial.

Edge computing deploys computing resources nearer the edge of the network instead of having them centralized. However, the edge's location depends on your vantage point. For our analysis, we'll use the diagram on the next page to classify edge locations.

For the purposes of this report, edge locations of relevance include:

- **Enterprise or customer edge** — This refers to computing, storage, and networking resources located on-premises at customer locations. These locations could be factories, office buildings, mines, oilfields, or semi-public venues like airports, stadiums, theme parks, and shipping ports.
- **Cloud or telco edge** — These are off-premises locations hosted by either public cloud providers or telecom operators. Topologically close to enterprise locations, these include mobile switching centers (MSCs), wireline central offices (COs), and mobile cell sites. Additionally, they may be localized public cloud edges placed near high-density metropolitan areas such as Los Angeles, Chicago, New York, London, Frankfurt, Seoul, Tokyo, or Toronto.

Acknowledgements:

We thank the following companies for providing their perspective for this report: Harman, HERE, and Savari.

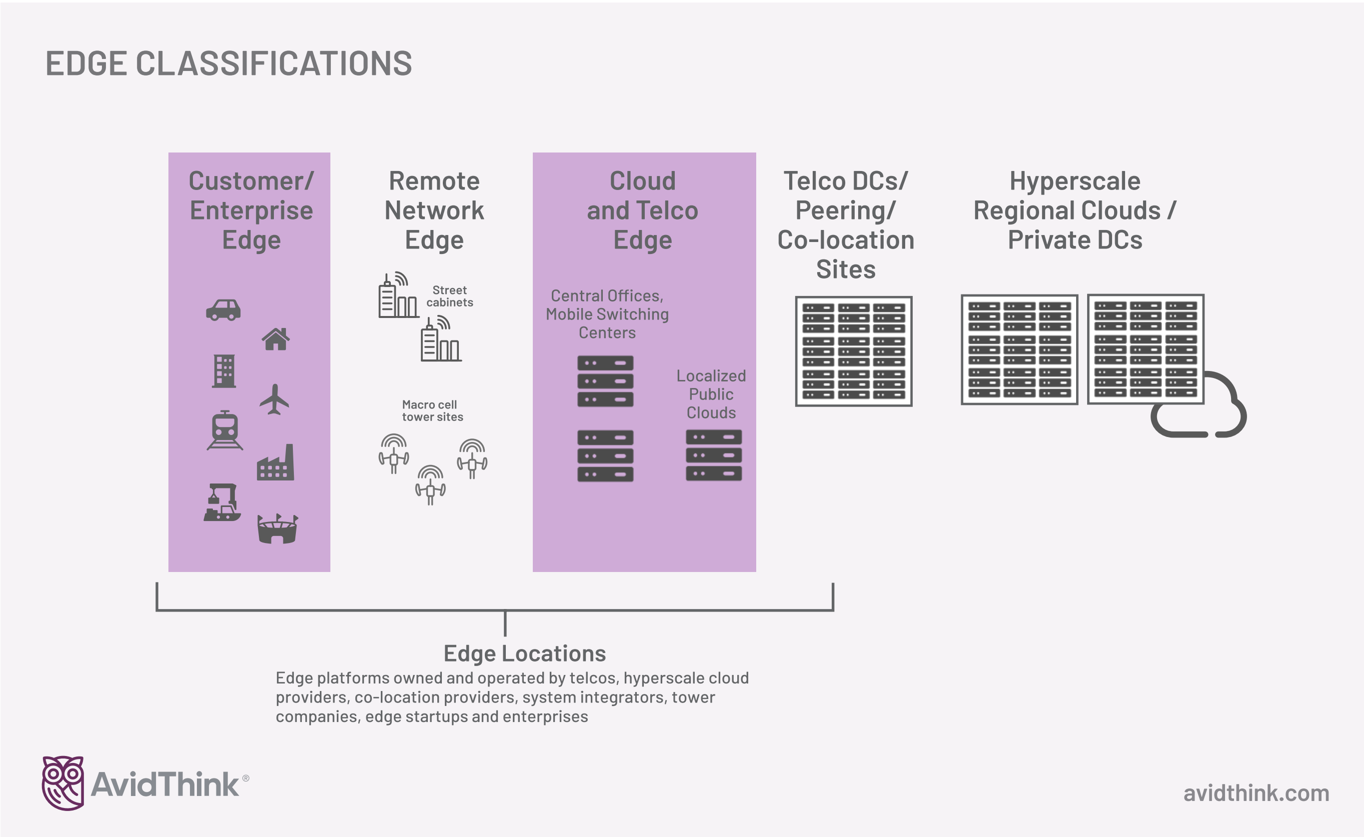

With the shift towards 5G-powered mobility, one edge in particular has garnered significant interest — the mobile edge. The telecommunications industry has been looking into embedding computing resources within the mobile radio network to improve user experience and application performance. Originally called mobile edge compute, the acronym MEC was redefined as multiple-access edge compute to broaden its applicability, and the European Telecommunications Standards Institute (ETSI) has led efforts to promote it.

MEC embedded deep within the radio access network (RAN) is still in early trials. However, edge computing resources located upstream in the MSCs or mobile service access points (SAP), where traffic from multiple RANs are aggregated, are now available for public access. Hyperscale cloud providers such as Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure are partnering with mobile network operators (MNOs) globally to extend their cloud computing platforms into the mobile network. As of this research brief's publication, AWS Wavelength service is the only generally-available public cloud service embedded in a telco network. AWS Wavelength's footprint has continued to grow since its launch in August 2020, and the service is available in multiple dozens of cities across the US, Canada, Europe, and Asia. Carrier partners for AWS Wavelength include Verizon, Vodafone, KDDI, SKT, and Bell Canada.

Benefits of Mobile Edge Computing

For the remainder of this brief, when we discuss mobile edge computing, we'll focus on the class of solutions represented by AWS Wavelength and similar potential offerings from other hyperscale cloud providers. We will reserve discussions around other edge locations to future reports.

As it stands, MSC-based edge computing platforms offer significant benefits, including:

- **Lower latency** — By having computing and storage resources located at or near MSCs/MSOs, applications and data hosted in these locations can be served directly to mobile devices, including vehicles, without incurring additional transmission latency or the vagaries of internet routing.
- **Reduced jitter with increased reliability** — Due to the reduction in hops and devices that the packets have to traverse, connectivity to edge sites is more reliable and suffers less jitter.
- **Higher bandwidth** — Depending on the uplink capacity from the MSCs to the rest of the network, heavy mobile traffic can bottleneck in backhaul network connections. But data stored and generated at the mobile edge and transmitted to mobile devices does not clog those backhaul connections. In this way, the mobile edge could act as an intermediary, caching data both ways, optimizing backhaul links, and scheduling transfers at off-peak times.
- **Backhaul cost reduction** — As a corollary to the previous benefit, if data is generated and processed at the edge, then backhaul links can be sized appropriately, resulting in cost savings. For use cases involving preprocessing and filtering of data before uploading to a central location, bandwidth reduction can lower costs.

One of the more holistic benefits of mobile edge clouds is the ability to spin up nearby cloud computing resources on demand.

- **Convenience of on-demand cloud services** — One of the more holistic benefits of mobile edge clouds is the ability to spin up nearby cloud computing resources on demand. Enterprises could build their own edges by purchasing racks of equipment and placing them in data centers near- or on-premises. However, that approach lacks the elasticity and on-demand benefit of edge clouds hosted by carriers or cloud providers. It also necessitates upfront capital expenditure instead of a pay-as-you-go billing model. The on-demand nature of the edge can prove valuable for emergency operations during natural disasters or handling data needs on crowded freeways during rush hour.

MSC-based mobile edges provide a good balance between scale and proximity. While more limited than a regional cloud data center, MSCs or other mobile aggregation sites have a larger scale than cell tower locations while providing low latencies.

Use cases that benefit from these mobile edges include video analytics, location services, IoT applications, augmented reality/virtual reality (AR/VR), content distribution, and caching. For the automotive industry, early use cases being explored include safety, control, entertainment, and information services. Another use case for automobile manufacturers is factory and industrial automation, but we'll cover that in another research brief on Industry 4.0.

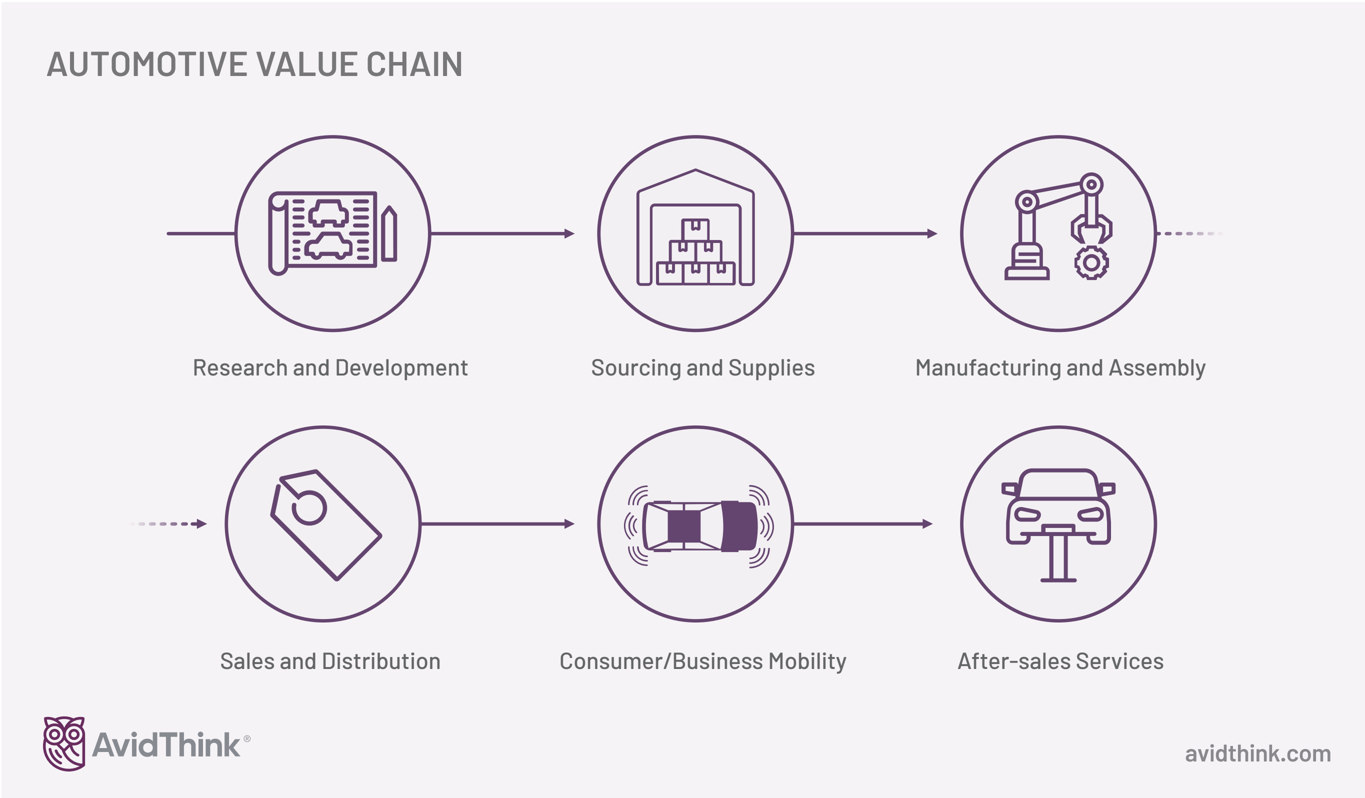

Trends in the Automotive Industry

The automotive sector is often showcased in the journey to 5G. Notably, autonomous driving — controlled and assisted by edge computing — has been touted as one of the significant use cases. However, fully autonomous vehicles dependent on edge infrastructure face regulatory approval and will require ubiquitous and reliable 5G coverage, which will take time to rollout Nevertheless, autonomous vehicles with onboard intelligence will continue to evolve, and the mobile edge will play a significant role in the path leading to fully autonomous (known as level 5 autonomy) in the commonly accepted framework from SAE International.

In any case, the whole automotive value chain (above), from R&D to manufacturing to after-sales service, is undergoing transformation driven by societal, business and technology changes. **Autonomous** vehicles is one of the four pillars of change that car industry pundits are converging on. The remaining three pillars are **connected**, **electrification** and **sharing** (collectively known as ACES). We'll delve into these pillars below.

Electrification

With increasing regulatory pressure from governments and growing customer acceptance, electric vehicles (including hybrid, plug-in, battery, and even fuel cell) comprise a growing share of the global automotive market. Tax incentives and other benefits — carpool lane access, parking discounts, discounted charging — offered by state and federal governments will continue to drive the adoption of these vehicles. Further improvements in battery and electric drivetrain technologies, significant drops in battery costs per KWH, and build out of charging infrastructure has accelerated consumer demand.

As vehicles transform from combustion engines into electrified systems with sophisticated sensor-driven systems, a move into software-based control is accelerated. The process of electrification can alter the vehicle's drivetrain and system architecture. Modern vehicles today have numerous electronic control units (ECU) and telematic control units (TCU) that make up the electrical control system on a car. Increasingly, many of these will be software controlled. The automotive industry has adopted the "software-driven" moniker, and electrification is a fundamental first step to the software-driven car.

Electrification allows sophisticated computing, in the form of AI and machine learning, to process the input from onboard sensors to promote efficiency and aid safety. Outfitted with digital signal processing chips, AI chips, sizable flash-based storage, and numerous mobile antennas, today's vehicles look like mini data centers on wheels, a far cry from yesterday's gasoline-powered cars. Electric systems lay a foundation for connected and autonomous vehicles and are driving the re-architecture of the car into a series of modular subsystems that can be tied together via a software platform. The transformation of the car is mirroring the same process in data centers and telecommunication networks — modular systems, disaggregation, commodity hardware, and software-driven intelligence. Eventually, we will see the evolution of the onboard operating system, applications, APIs, and semblance of virtualization.

When combined with connectivity, the electrified vehicle is propelled into the realm of a software-upgradeable computing platform.

Connected Cars

The electrified car provides a platform upon which other capabilities can be built. When combined with connectivity, the vehicle is propelled into the realm of a software-upgradeable computing platform that mimics a mobile phone. Yet it also ties into the different physical capabilities of the vehicle, including sensors, motors, brakes, suspension, multiple displays, sound systems, and climate control systems.

Connectivity provides the opportunity for media streaming, gaming, and other services during transit. It also powers information consumption, such as high-definition maps, road condition warnings, and other safety features that go beyond those provided by a mobile phone.

The connected vehicle provides manufacturers and third parties the opportunity to upstream critical sensor and performance information. Smart and predictive maintenance can be scheduled based on deep analytics or use of AI. Performance data can be analyzed and software upgrades designed to improve the handling and safety of vehicles. Connected vehicles can even have their software upgraded over-the-air (OTA), like today's mobile devices.

Select car manufacturers view the connected vehicle as a docking platform for personal mobile devices, enhancing the consumer experience by leveraging the compute and mobile capabilities of phones while extending to include vehicle resources. Other manufacturers are banking on the car as the primary connectivity device, with a sweeping digital cockpit that goes beyond the simple interactivity of a mobile phone. Regardless, all envision the vehicle evolving into a mobile entertainment and informational hub, especially when paired with autonomous operation.

Autonomous Driving and Safety

The move towards electrification and connected vehicles will also enable autonomous vehicle operation. Prior to autonomous driving, vehicles sported improved safety and convenience features under the umbrella of advanced driver-assistance systems (ADAS). Assisted parking, lane changing, lane departure warnings, collision avoidance, and other features have helped lay the groundwork for consumers, regulators, and enterprises to embrace the fully autonomous vehicle.

There's little doubt that fully autonomous driving systems will capture the attention of the public and lawmakers over the next few years as we progress towards SAE level 5 full autonomy. Whether used for public transportation (autonomous buses), ride sharing, private vehicles, commercial deliveries, or trucks shipping goods from ports to warehouses, autonomous vehicles will make headlines, both positive and negative.

Autonomous vehicle success will initially be dependent on the sophistication and accuracy of onboard systems — those monitoring, tracking, and reacting to hundreds of megabytes and terabytes of information generated daily. Imagine the size of data streams from onboard radar, lidar, high-resolution near-field and far-field cameras, microphones, and inertial sensors. Some will be processed onboard, but substantial data will be offloaded to the edge and cloud. Early estimates suggest that an autonomous vehicle will generate and consume up to 40TB of data for every eight hours of driving — generating a high-definition map of a city will require 1.5TB alone — without the ongoing dynamic updates and downloads. Not all of that traffic needs to be offloaded over a mobile network; some can be transferred during charging or while parked, but to backhaul even a portion will put a significant load on networks.

Edge computing and cloud computing will play an essential role in assisting with the analysis of these large data sets and generating models and actions that improve the comfort and safety of autonomous or assisted operations.

Sharing and Subscription

Consumer behavior around vehicles has rapidly changed in the past few years. Ride sharing and vehicle sharing have dominated the landscape across major metro areas worldwide. Vehicle manufacturers like Volvo, Porsche, Mercedes-Benz, Audi, BMW, and Nissan are market testing vehicle subscription services (different from leasing) that wrap all aspects of vehicle ownership into a service — maintenance, insurance, etc. There is indication of an eventual market for shared mobility, where vehicles can be treated like an AirBnB lease and rented for a limited duration.

Mobility as a service aligns with the autonomous capabilities being developed today, and the connected car enables ongoing tracking and constant communication with the vehicle. The continual monitoring will support advanced enterprise fleet management features and help car owners or manufacturers who temporarily rent out their vehicles for use.

Likewise, connected vehicles can have safeguards that prevent vehicles from being driven out of a fenced geographic area at high speeds or dangerous maneuvers. Real-time alerts allow vehicle owners to be notified in a timely manner to mitigate and address any violations.

INTERNAL AND EXTERNAL FACTORS IMPACTING AUTOMOBILES

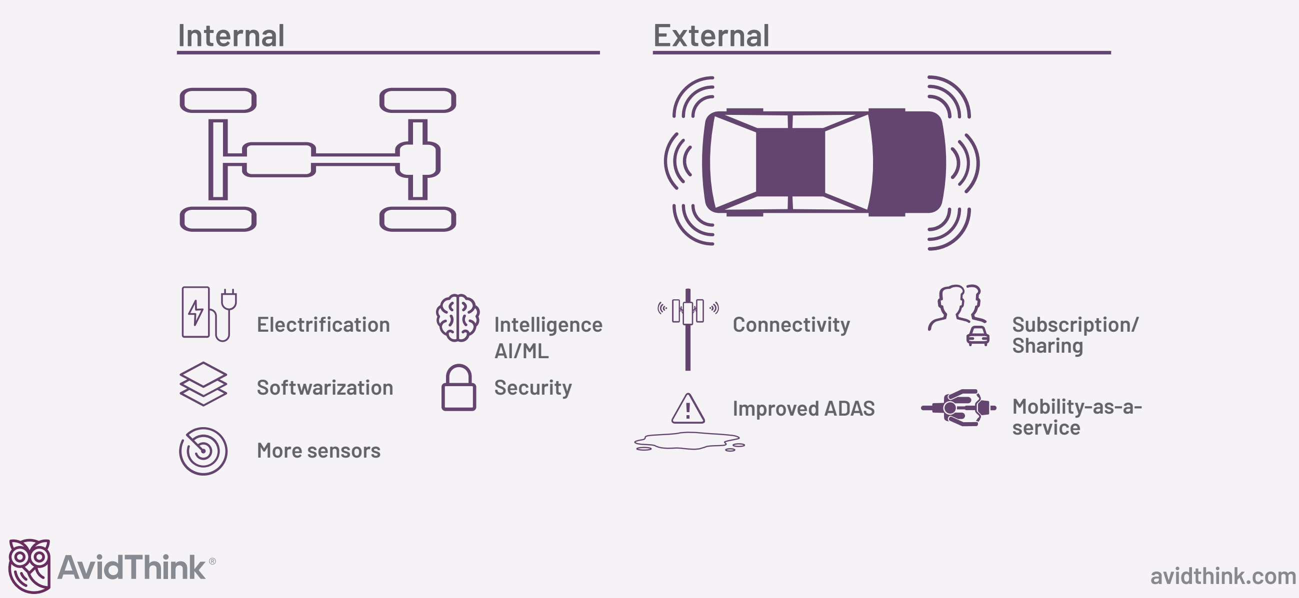

Other Considerations

Beyond the core ACES pillars, there are other trends that impact the automotive ecosystem and value chain. We'll start with a major consumer trend that's emerging in major cities across the world.

Multimodal Transportation

Consumers are increasingly adopting multimodal transport that goes beyond driving point-to-point in a vehicle. Using a combination of cars, ride sharing, subways, trains, bicycles, and scooters, consumers are able to leverage the connectedness of today's transportation options to plan and optimize their travel. Likewise, consumers are looking for delivery of goods to where they work and live, instead of making trips to retail centers.

This new era of transportation could prompt a decline in the private vehicle market. Conversely, as the convenience of ride sharing drives up the number of vehicle trips in major cities, increased utilization and wear and tear could result in lower lifespans and increase vehicle purchases. At the same time, we expect to see diversification in vehicles, with new models specifically built for ride sharing that are more robust, able to handle high mileage, and with specific features for passenger compartments like isolated entertainment and audio for privacy. Fleet features for tracking and proactive maintenance will likely be built into vehicles targeted at business use.

Dynamic Changes in Automotive Ecosystem

The electrification of the vehicle and evolution of the drivetrain, coupled with the increased move to software-centricity, opens the market to new competition. Ride-sharing companies like Uber and technology and cloud companies like Google, Amazon, and Apple represent potential partners but also competition. New entrants with different skill sets like Tesla also represent a threat to the incumbents like GM, Ford, and BMW.

The electrification of the vehicle and evolution of the drivetrain, coupled with the increased move to software-centricity, opens the market to new competition. Ride-sharing companies like Uber and technology and cloud companies like Google, Amazon, and Apple represent potential partners but also competition.

At the same time, the software focus and the vehicle-as-a-platform theme encourages new startups in the field who are treating the car like a mobile phone. New applications for safety and ADAS, infotainment, communications, optimization, fleet management, and maintenance will show up. Some startups will interact directly with consumers, while others will work closely with vehicle manufacturers or other parts of the automobile supply chain.

Importance of Security

As with any connected and software-centric platform, security needs to be a key consideration. Some connected vehicle platforms today feature a security module that acts as a gatekeeper for communications between the ECUs across different vehicle zones, watching for malicious actors. The same gatekeeper is used to inspect bidirectional communications between the vehicle and the network. Given the potential for injury and loss of life from an errant vehicle, vehicle manufacturers understand the critical nature of preventing platform compromise.

Automotive Use Cases for the Mobile Edge

Mobile 5G ties in closely with the evolution of the connected vehicle. For instance, a group of automotive, technology, and telecommunications service provider companies have come together to form the 5G Automotive Association (5GAA). A member of the 3GPP (the organization focused on setting the standards for 5G), 5GAA is focused on developing mobile use cases for the transportation industry. Four years after its founding, the 5GAA has grown to over 120 members. An alternate association focused on the use of edge for automobiles is the 5G Automotive Edge Computing Consortium (AECC) led by Toyota, Ericsson, NTT, Cisco, and Dell EMC. At the same time, ETSI, a telecommunications standards body focused on standardizing MEC, is collaborating with these other entities to further automotive MEC frameworks.

The work done by these groups involves mobile network operators (MNOs) embedding capabilities within the mobile core to facilitate edge service discovery, vehicle mobility handling, and fast packet routing for the automotive use cases. Furthermore, some mobile edge architectures are dependent on the presence of road-side units (RSUs): compute units located in proximity to roads and highways that provide computing support with ultra-low-latency access from vehicles. Due to the costs and logistics required to deploy these RSUs, limited sections of roads globally have them installed for early testing. Some in the industry postulate that 5G network radios, especially small cells for mmWave deployments in denser urban areas, might be co-located with these RSUs, or eventually might play the role of RSUs.

In any case, many MNOs today are focused on the basic rollout of 5G. They have started converting their 4G radios to 5G New Radios (NR) to benefit from spectrum efficiencies and capacity while keeping their existing 4G packet cores for non-standalone (5G NSA) initial deployments. More advanced MNOs are migrating their packet cores to next-generation 5G cores (5GC) to reap the benefits of a modern, more scalable architecture.

5GAA and AECC have begun sharing recommendations on mobile edge architecture, indicating which parts of a 4G or 5G core need to be involved in facilitating connected vehicle communication. There's work needed to standardize edge data offload, edge server selection, vehicle reachability, access network selection, provisioning, and configuration updates, as well as facilitating opportunistic data transfer without overloading the network.

While this work continues, MNOs and parts of the automotive ecosystem are already trialing infotainment, safety, and autonomous use cases using proprietary edge backends and regional cloud-based services. MNOs today, given the maturity of their 5G infrastructure, have not yet prioritized making changes to support automotive edge computing explicitly and certainly not in a standardized way.

5GAA and C-V2X

5GAA is looking to bring to market use cases based on Cellular-V2X (C-V2X), where V2X represents vehicle-to-any-thing communication. The "anything" could be infrastructure (V2I) like traffic lights, another vehicle (V2V), the network (V2N), pedestrians (V2P), or other mobile devices. The C-V2X framework is highly dependent on advances in the cellular network and is a single framework that covers both direct or sidelink connections (V2V, V2I, V2P) through the PC5 interface and network or uplink/downlink connections (V2N) through the Uu interface.

The LTE portion of C-V2X was part of the 3GPP Release 14 and 15 specifications (on which mobile equipment providers and carriers build out their mobile networks). In comparison, the 5G portion, called NR-V2X, was part of Release 16 and complemented LTE-V2X to support automated driving.

As the MNOs and automotive organizations continue to hash out changes needed for the radio access network (RAN) infrastructure, we are excited by the availability of the new wave of mobile edge that's situated close to the RAN. Services like AWS Wavelength (or those to be launched by the other hyperscalers) provide a RAN-agnostic solution that can host portions of automotive applications that need proximity to the vehicles on the road.

These mobile edges can immediately provide improvements in performance, think latency, jitter, throughput, and reliability. As the RAN matures to better support automotive use cases, automotive applications can further evolve to take advantage of those additional capabilities. In the meantime, we see early adoption by both vehicle manufacturers and automotive application developers in using these types of mobile edges. We'll delve into these use cases as tied to the ACES trend. In particular, the electrification and connectivity initiatives enable vehicles to connect in real time to the edge, and the low latencies and compute capabilities provide important assistance for autonomous vehicle operations.

Infotainment

- **Streaming content caches** — For in-car entertainment, keeping popular content on the mobile edge as part of a content delivery network (CDN) to allow fast access can improve the consumer experience in today's connected vehicles.
- **High-definition maps** — High-definition map tiles or other localized information can be stored at the mobile edge, reducing the access time to pull high-definition content and allowing for high interactivity with rich content. Some trial implementation includes the use of augmented reality (AR) to overlay critical information over real-world views through the windscreen. Due to the local nature of the data, storing voluminous HD content at the edge reduces network backhaul use while maximizing user experience.

Safety

- **Real-time situational awareness** — Local urgent events, like icy roads or mudslides, can be updated and relayed with minimal delay to ensure that up to the second or even millisecond updates are available to all vehicles in the vicinity with high reliability — this is where regional clouds don't have sufficient proximity to assist. Any hard braking or collisions ahead can be transmitted with ultra-low latencies to approaching drivers, allowing them to react, or in the case of assisted or autonomous, allow the vehicle to take preventive or evasive action.
- **See-through capability** — The mobile edge can also participate in "see-through" operations that allow vehicles to understand what's happening on the road ahead, even if their view is blocked by large trucks. High-resolution video feeds from vehicles ahead can be transmitted in near real-time and displayed to vehicles behind, providing visibility even without line-of-sight.
- **Vulnerable road user detection, intersection movement assist** — Vulnerable road users (VRU) like pedestrians or cyclists upfront can be detected via cameras with AI inference assistance from the mobile edge. Even rear cross-traffic

evaluation can be handled more comprehensively than with just onboard information and sensors. This information at the edge can be rapidly transmitted to vehicles in the vicinity, particularly those without direct visibility of upcoming intersections or vehicles turning the corner.

Convenience and Comfort

- **Road conditions adjustment** — Local road conditions like potholes, large pools of water, or upcoming high-fidelity information about the route can be transmitted to vehicles to inform the driver which lane to stay in, or to update the traction control systems, or to modify the suspension algorithms and improve overall vehicle handling.
- **Vehicle software updates (OTA)** — Software updates will usually happen when the vehicle is parked, and latency is less of an issue. However, the mobile edge acts as a software update cache, reducing the load on backhaul networks and speeding up the update process.

Driving Assistance and Autonomy

- **Autonomous driving assistance** — While the logic and sensors necessary for fully autonomous will likely be onboard cars for the foreseeable future, the mobile edge can provide an enhancement layer of information that contains localized road conditions, traffic patterns, and any specific elements that might improve the autonomous vehicle logic. Real-time sensor feeds from vehicles ahead can be processed at the edge and provided to ensure better route management or safer driving. The other role that the mobile edge can play is in precision location, going far beyond GPS to understand the location of the car on the roadway. This detailed location can be used for driving assistance and safety.
- **Traffic efficiency** — Localized traffic analysis can be performed and optimizations run to provide vehicles with improved routing to optimize traffic flow, avoiding slowdowns or traffic jams.
- **Cooperative lane change** — For lane changing, vehicles can leverage the mobile edge to signal other nearby vehicles in the target lane of their intention, ensuring a safe and smooth operation.
- **Automated or assistive parking** — The edge can act as a localized orchestrator for automated parking operations in a local area or a parking lot, providing the necessary intelligence and information to autonomously park a vehicle or at least provide guidance to the vehicle.

Security

- **Malware and attack prevention** — The mobile edge can host malware detection and network protection that scrubs traffic being downloaded to the vehicles. Likewise, it can monitor uplink traffic to detect if a vehicle has been compromised and prevent the spread of malware or attacks to other vehicles. It can protect against denial-of-service attacks both ways — attacks that attempt to compromise or overload the V2X infrastructure.
- **Identity and authentication** — For many V2X functions, ascertaining the identity of vehicles and the infrastructure is critical to ensure that only trusted entities are in communication over the network. The mobile edge can perform the necessary authentication and authorization functions with minimal latency, ensuring that security does not get in the way of critical real-time functions.

Mobile Edge Benefits Summary for the Automotive Sector

With the large number of mobile edge C-V2X trials going on worldwide (East Asia, Europe, and North America), we expect innovative use cases to emerge. Regardless, we can summarize the benefits of a mobile edge platform for the new generation of automobiles as follows:

- **Lower latency** — While not the sole benefit, it is a primary driver behind many use cases. We expect to see consistent sub-20ms or sub-10ms for latency-critical use cases on the public mobile network. These low latencies allow the use of edge computing for safety and autonomous driving assistance when available. Likewise, low latencies can enhance the user infotainment experience, reducing wait times and allowing for responsive interaction.
- **Reliable and consistent network** — The reduced number of network elements between the mobile device, application servers, and data reduce variability and packet loss. The increased reliability makes the mobile edge a strong candidate to host critical information for safety and autonomous operations, versus regional clouds, which are likely to suffer from higher packet loss rates or routing inconsistencies. This is especially true for safety-related use cases described above: vulnerable road user detection, real-time road conditions, and hazard warnings, the mobile edge has the reliability — and low latency — needed.
- **Backhaul traffic reduction** — Vehicles are expected to generate a significant amount of data in safety and infotainment use cases. To the extent that a significant portion of this data is localized, using the mobile edge can reduce the amount of traffic directed to regional clouds, keeping the backhaul network lightly-loaded.
- **Reduced learning curve** — For a mobile edge network that's paired with a public cloud, one of the most underestimated attributes is the frictionless developer experience in extending applications to the edge. Automotive applications already built on cloud infrastructure can simply be extended to the edge without scaling another learning curve. Application developers won't need to worry about how each mobile operator will implement their edge cloud; instead, they will be able to reuse their development methodology, application monitoring constructs, and infrastructure. This silo-free approach speeds up innovation at the edge.

Automotive applications already built on cloud infrastructure can simply be extended to the edge without scaling another learning curve. Application developers won't need to worry about how each mobile operator will implement their edge cloud; instead, they will be able to reuse their development methodology, application monitoring constructs, and infrastructure.

- **Edge intelligence** — Edge clouds have AI and GPU acceleration that can be used to offload onboard processors. By leveraging edge systems, automobile manufacturers can provide richer applications without compromising space, weight, or power in the vehicle.
- **Improved data jurisdiction compliance** — Mobile edge clouds will tend to be located in country. For regions such as the E.U., where regional clouds might be located in a different country, utilizing edge clouds for data processing of sensitive data will aid compliance with privacy and data processing regulations.
- **Reduced time to market worldwide** — The full frameworks described by the 3GPP, 5GAA, AECC, and ETSI will take time to standardize and be implemented by MNOs worldwide. However, using the mobile edge at MSCs, automotive application developers can immediately benefit from the improved network performance and available computing and

storage present. As the standards bodies and MNOs roll out more sophisticated capabilities in network, vehicles can evolve to incorporate them. Using a mobile edge cloud offering that's partnered with a hyperscaler allows automotive companies to achieve a global footprint with a quicker time to market. With today's automotive companies being global, or at least multiregional, relying on widely available infrastructure allows a uniform strategy to serve vehicles worldwide, without waiting for individual negotiations with MNOs in different markets.

Fundamentally, the mobile edge provides automotive companies with the reliability, agility, and consistency they need to meet the consumer demands. This edge provides a uniform platform on which to run automotive applications with low latency and high reliability, reducing the need to navigate every local mobile market's complexities.

Mobile Edge Considerations for the Automotive Sector

While the mobile edge comes with a host of potential benefits, there are considerations that companies in the automotive ecosystem should be aware of. These include:

- **Edge platforms are evolving** — The mobile edge is an early technology. While it uses mature and trusted cloud platforms, the specific placement of these resources is a new operation, as is the tighter coupling with mobile networks. Since the automotive sector works on vehicle designs with a two-to-three-year lead time, platform architects need to factor in some level of flexibility to track the rapid evolution of the mobile edge platform. We should also note that mobile networks will be evolving to conform to 3GPP support for C-V2X, and the 5GAA, AECC, and other organizations will push to introduce new MNO-impacting frameworks. Developers will want to track the changes in this space closely. Nevertheless, the advantage of leveraging this hyperscaler-MNO partnered mobile edge is that the approach is agnostic to RAN changes and will benefit from any underlying improvements to handle the vehicle mobility.
- **Financial tradeoffs exist** — Some workloads do not need the low latency or reliable connections at the edge. For instance, regional clouds may make sense for a vast majority of infotainment content and will cost 20-30% less than the edge for compute and storage resources. Likewise, for non-time-critical data offload, data transfer off-hours when vehicles are garaged or at the depot, or even during charging might make more sense. Mobile data costs are always going to be higher than direct wireline connectivity or private wireless networks.
- **Mobile edge coverage is nascent** — Publicly available mobile edge services are starting to roll out. The number of regions served is limited today but multiplying. For instance, AWS Wavelength has launched across North America, in Europe, and in Asia. Yet Wavelength Zones in an MNO network only assures low-latency access for that particular mobile operator's customers. Nevertheless, more rollouts are planned across all hyperscalers in partnership with global MNOs, which will mitigate this concern in the longer term.
- **Application orchestration across edges will add complexity** — Related to edge platform maturity, we need to consider the orchestration of application components across multiple locations and optimization to maximize performance and user experience while minimizing costs and adhering to business-set constraints. Some of these considerations may be dynamic and involve real-time user load (e.g., as vehicles move across locales) or density of users in a region at a time (e.g., traffic jams). While the hyperscale clouds may provide underlying infrastructure management across regional clouds and mobile edge zones, application developers will need to carefully consider workload placement logic and leverage automation provided by the underlying cloud platforms.

As with any promising new technology, there will be bumps in the road in early deployments. However, given the close intertwining of 3GPP 5G standards with mobile use cases, there's undoubtedly a significant opportunity in this arena.

Getting Started with the Mobile Edge

The connected car and autonomous vehicle market are nascent yet burgeoning, and what's clear from our conversations with innovative players in this space is that there is much to learn. Many have shared that real-world trials and deployments are the only way to understand how to evolve and address the needs stemming from the ACES trends.

While mobile networks evolve their RANs to support vehicle mobility, connectedness, and autonomy, the advent of the hyperscaler-partnered mobile edge helps accelerate the ecosystem. Automotive companies can start running public trials and rolling out applications that leverage low latencies and higher connection reliability.

For automotive companies at the forefront, AvidThink recommends taking the following steps:

- **Architect automotive applications for a distributed cloud** — The computing platform should be viewed as a continuum. Core applications can be hosted in regional clouds that have the compute and cost-effectiveness necessary for scale. Other latency-sensitive components or components that need proximity to the vehicle should be built to run across a continuum of clouds (multiple edges) and even on-vehicle. Unlike other use cases where the end-user device is constrained, today's vehicles look like micro data centers on wheels and can be leveraged as part of the distributed cloud.
- **Gracefully handle disconnected operations** — Recognize that the network should be viewed as an enhancement layer for applications. While mobile networks are more ubiquitous today, there will be situations of no or inadequate coverage. The applications should be able to defer and queue operations or run in a degraded local mode without violating safety guidelines.
- **Gain learnings early** — The automotive design lifecycle means that any understandings and design changes will take at least a year or two before they get into production models. Therefore, it's critical to obtain learnings as early as possible to ensure the most robust and innovative designs are incorporated into future model years.

To get started, companies can begin trials on commercial edge offerings available today. The mobile edge will play a critical part in the future of mobility and automotive. Driving early engagement will put automotive companies on the road to success.

AvidThink, LLC
1900 Camden Ave
San Jose, California 95124 USA
avidthink.com