WHITE PAPER

Automotive Ethernet: Enabling the Future of Autonomous Driving

KEYSIGHT

Introduction

Major innovation drivers in the automotive industry span three categories: enhanced safety, a cleaner environment, and improved convenience with connectivity. To achieve these goals, automakers, automotive suppliers, governments, academia, and even nontraditional automotive players such as wireless chipset makers, mobile device makers, and wireless service providers are developing advanced driver-assistance systems (ADAS), connected car technologies, and autonomous vehicles.

ADAS and autonomous vehicles require a high-bandwidth and low-latency network to connect sensors, cameras, diagnostics, communications, and central artificial intelligence (AI).

The wiring harness is the third heaviest component in a vehicle and the third most costly system. Wiring harness installation represents 50% of labor costs during automobile assembly.

Automotive Ethernet is the emerging solution to these challenges in the same way that Wi-Fi is the foundation for dedicated short-range communications. Ethernet is a well-known, trusted, and ubiquitous solution in traditional local area networking (LAN). The advantages of Ethernet — multipoint connections, higher bandwidth, and low latency — are attractive to automobile manufacturers. However, traditional Ethernet is too noisy and interference-sensitive for use in automobiles. The Institute of Electrical and Electronics Engineers (IEEE) has new standards and protocols to deliver on the specific needs of the auto industry.

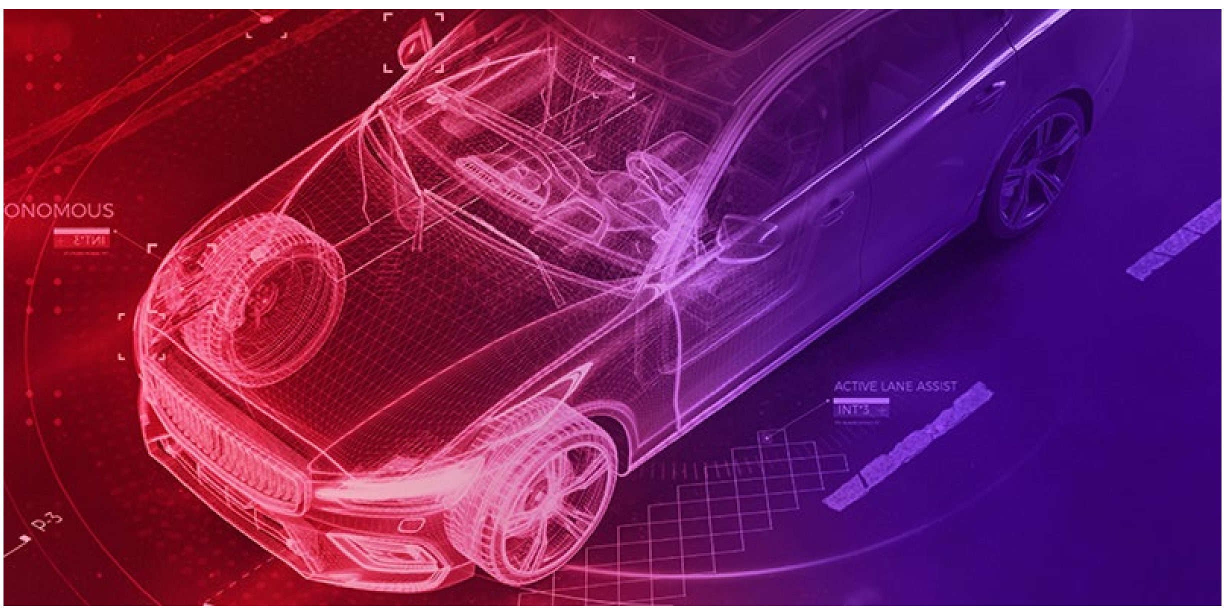

Autonomous Driving Equals More Data at Faster Speed

Technologies that enable autonomous vehicles span an array of electronic components. The first category allows sensor fusions across radar, light detection and ranging (lidar), and cameras. The second category covers wireless communications for vehicle-to-vehicle, vehicle-to-network, vehicle-to-infrastructure, vehicle-to-pedestrian, vehicle-to-utility, and eventually vehicle-to-everything (V2X) systems. Adjacent elements, such as high-definition mapping with high-precision navigation systems, powerful signal processing, and AI, round out the required components for autonomous driving.

These technologies generate, send, receive, store, and process enormous amounts of data. For example, a lidar module provides highly accurate, high-resolution 3D and 360-degree imaging data around the car. A lidar module can generate 70 Mbps. A camera can generate 40 Mbps, 100 Kbps for a radar module, and 50 Kbps for a navigation system.

Autonomous driving

- Autonomous vehicles and advanced driver-assistance systems require higher bandwidth and lower latency.
- Automotive Ethernet has become the new backbone for faster automotive networks.
- Comprehensive testing of the transmitter, receiver, link segment, and higher-layer protocol functions ensures successful implementation.

Moreover, with higher levels of autonomous driving systems, the number of individual sensors will also dramatically increase, thereby increasing total data generation. For example, a Level 2 autonomous driving system provides longitudinal and transverse guidance, so drivers can free their hands and temporarily avert their eyes. It may use five radar sensors and five cameras. A fully autonomous driving system (Levels 4 and 5) would require up to 20 radar sensors, six cameras, and V2X communications. Most experts agree that an autonomous vehicle will typically generate anywhere from 5TB and 20TB of data daily. The vehicle must transmit, store, and share this data with very short latency on a reliable high- speed network, building the case for a high-throughput, low-latency network based on automotive Ethernet.

Automotive serial buses

We can better understand why autonomous vehicles and ADAS cameras and sensors require automotive Ethernet by reviewing the traditional automotive serial buses, including Controller Area Network (CAN), low-voltage differential signaling (LVDS), Local Interconnect Network (LIN), Media Oriented Systems Transport (MOST), FlexRay, and Controller Area Network Flexible Data-Rate (CAN FD).

CAN — 1983

CAN, developed by Bosch, is a shared serial bus that runs at up to 1 Mbps. The International Organization for Standardization (ISO) published the CAN standard in 1993. Its advantages are cost-effectiveness and reliability. Its disadvantages are shared access and low bandwidth. Power trains, chassis, and body electronics use CAN. With the advent of advanced features like autonomous driving, the amount of data transmitted by the CAN bus is increasing exponentially. To enable these features, newer versions of the CAN protocol, like the CAN FD or CAN XL, appear on the market.

LVDS — 1994

LVDS is a high-speed, long-distance digital interface over two copper wires commonly used for high-speed image transmissions in a vehicle. It has a much lower cost than MOST, and many automakers use it for camera and video data. However, each LVDS link can interface with only one camera or video output at a time. There are new evolutions of automotive Serdes, both proprietary (GMSL, APiX) and standard (MIPI A-PHY and ASA), also being utilized by OEMs.

LIN — 1998

A consortium of automakers and technology partners developed LIN. It runs at 19,200 bits per second and requires only one shared wire, compared with two for CAN. LIN uses a master-slave architecture, while CAN treats all nodes as equal. LIN costs less than CAN, and its speed and cost suit body electronics such as mirrors, power seats, and accessories.

MOST — 1998

MOST is a high-speed multimedia network technology running up to 150 Mbps (MOST150) using fiber or copper interconnects. A MOST network can manage up to 64 MOST devices- in a ring configuration.

MOST has the advantage of relatively high bandwidth for the automotive market but is costly. It was initially intended only for camera or video connections.

FlexRay — 2000

FlexRay is a shared serial bus running at up to 10 Mbps. The FlexRay Consortium, a group of semiconductor manufacturers, automakers, and infrastructure providers, developed FlexRay. Unlike CAN, FlexRay has no built-in error recovery; the application layer handles error handling. It has a higher bandwidth than CAN but has a higher cost and shared media. High-performance power trains and timing critical safety-related systems, such as drive-by-wire, active suspension, and adaptive cruise control, use FlexRay. Automotive ethernet may replace FlexRay for future bandwidth-intensive, non-safety critical applications.

CAN FD — 2012

CAN FD, released by BOSCH in 2012, is an extension to the original CAN bus protocol intended to accommodate increases in bandwidth and payload support requirements in automotive networks. CAN FD increased the effective data rate up 8 Mbps and supports a flexible message payload up to 64 bytes per frame. CAN FD enable more accurate and near-real-time data by minimizing protocol delays and delivering higher bandwidth. CAN FD is backward compatible with existing CAN 2.0-based networks.

Automotive Ethernet — 2016

Although traditional automotive serial buses such as CAN, CAN FD, and FlexRay have played important roles in various automotive applications, it took a lot of work for those legacy buses to address the requirements set by sophisticated modern ECUs and complex in-vehicle networking. For example, most automotive serial buses cannot transmit at the 70 MB per second data rate required by a lidar or radar sensor. When integrating diverse sensing technologies and wireless communications, it is common to use lidar, radar, cameras, and V2X communications simultaneously. In this case, the amount of data transmitted is beyond the capacity of traditional automotive serial buses. This is one reason the auto industry is looking to automotive Ethernet to make autonomous driving and ADAS a reality.

CAN XL – 2022

CAN XL is the third generation CAN standard in development by the CAN in Automation (CiA) group that improves the well-established CAN FD protocol with increased payload data and faster bit rate. The CAN XL data phase speed is specified to reach 10 Mbps or more, depending on the transceiver capabilities and physical layer components. CAN XL has two modes of operation – fast mode and SIC mode (or slow mode).

What is automotive Ethernet?

Automotive Ethernet is a wired network that connects electronic components in a car. It is intended to meet the automotive industry's bandwidth, latency, synchronization, interference (for example, electromagnetic interference [EMI]), security, and network management requirements. Broadcom introduced the concept of automotive Ethernet, and the OPEN (One-Pair Ethernet) Alliance adopted and regulated it. Since then, IEEE has released standards for 100BASE-T1, 1000BASE-T1, 10BASE-T1S, MultiGBASE-T1, and most recently MultiGBASE-AU. Together, these standards define the general technology known as Automotive Ethernet.

The OPEN Alliance sponsored Broadcom's 100 Mbps BroadR-Reach as a multivendor licensed solution. The 100 Mbps PHY implementation uses technologies from 1 Gbps Ethernet to enable 100 Mbps transmission over a single pair in full duplex mode. This works with echo cancellation using more advanced encoding to reduce the base frequency to 66 MHz (from 125 MHz), allowing automotive Ethernet to meet automotive EMI specifications. IEEE and the OPEN Alliance have created and maintained physical-layer standards for 10 Mbps, 100 Mbps, 1000 Mbps, and Multigig bps automotive Ethernet in the IEEE 802.3 and 802.1 groups.

In its early years, Ethernet was intended for diagnostics and firmware updates through a single 100BASE- T1 1TPCE link from the DLC diagnostics port to the gateway. Figure 1 shows the evolving role of automotive Ethernet as the new backbone using higher-speed Gigabit Ethernet 1000BASE-T1 RTPGE links.

The 10BASE-T1(IEEE Std 802.3cg-2019) is a relatively new variation of the automotive Ethernet standard that offers a bandwidth of 10Mbps over a single pair of conductors. One of the variants specified under the IEEE is 10BASE-T1S, where S stands for short reach. A long-reach variant, called 10BASE-T1L, is defined for distances up to 1 km. The 10BASE-T1S uses a point-to-point or a multidrop topology where each node connects to a single cable. This eliminates the need for a switch, resulting in fewer cables as each cable only uses one pair of wires.

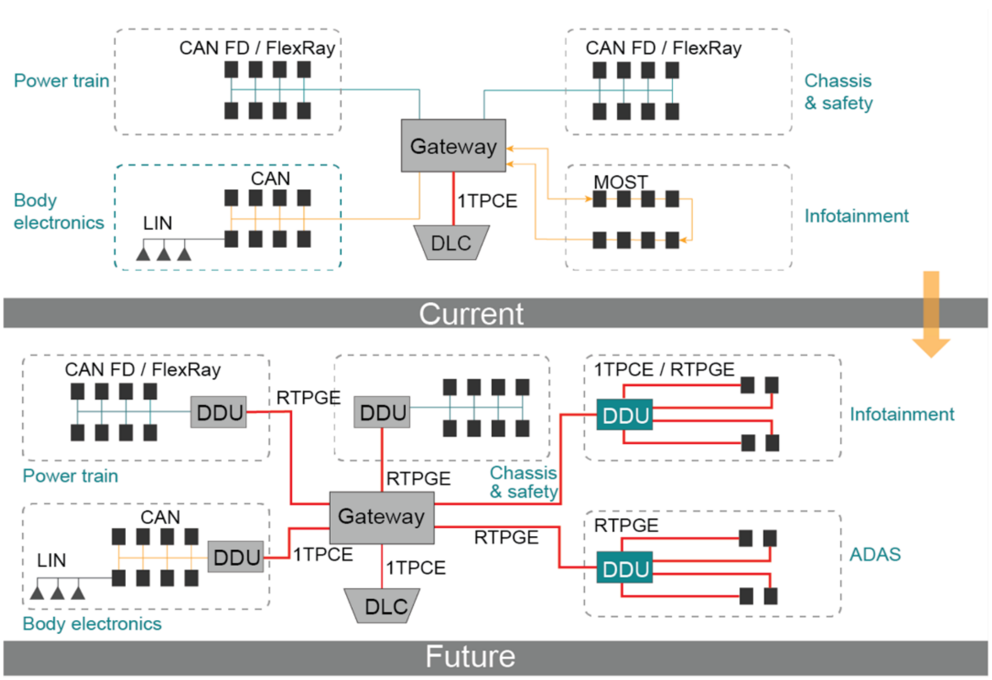

Figure 1.
The evolving role of automotive Ethernet

Why automotive Ethernet?

Automotive Ethernet is the obvious choice over traditional automotive serial buses for in-vehicle electronic systems connections and communications, autonomous driving, and ADAS. Car electronics architectures are becoming more complex, with more sensors, controllers, and interfaces, higher bandwidth requirements, more computers, and communication links.

The wiring harness that connects these systems is a car's third costliest and third heaviest component. Today, automakers use multiple proprietary standards for communication; most components use a dedicated wire or cable. Automotive Ethernet serves as the unifying standard for all communications. It is carried on a single cable pair from each electronic component to a central networking switch. Using unshielded twisted-pair cable and smaller compact connectors could reduce connectivity costs by up to 80% and cabling weight by up to 30%, according to a joint study by Broadcom and Bosch.

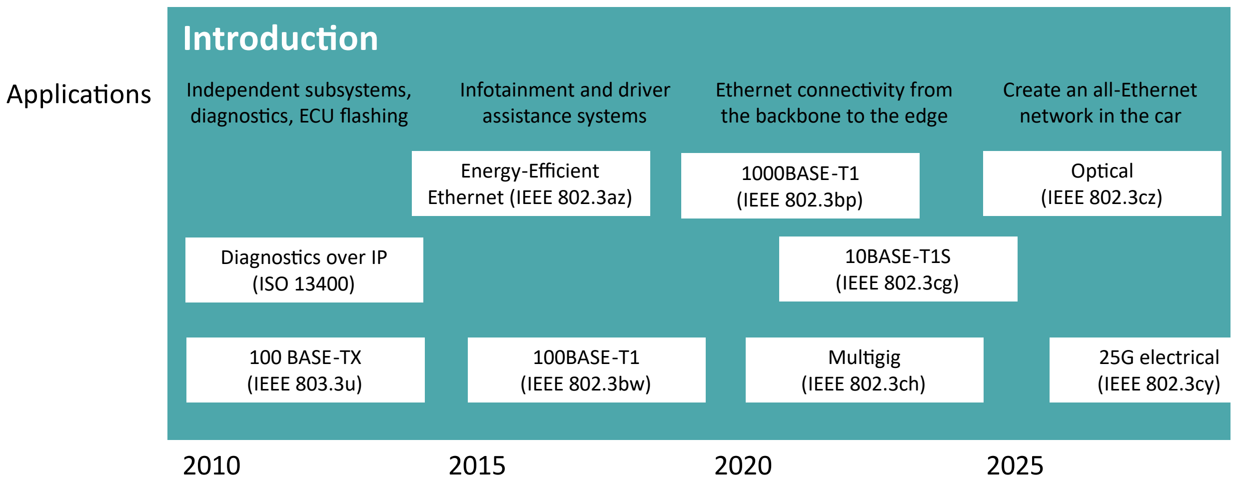

Figure 2. Automotive Ethernet progress

The evolution of automotive Ethernet technologies

Advances in multiple technologies areas have made the progress in Figure 2 possible.
Those advances include the following:

AUTOSAR (Automotive Open System Architecture)

AUTOSAR is an open and standardized automotive software architecture. Automobile manufacturers, suppliers, and toolmakers jointly develop the architecture. AUTOSAR includes the automotive TCP / IP stack used in automobiles. The auto industry has effectively agreed on AUTOSAR as the standard, with different automakers competing on implementation rather than on the standard itself. The standard implementation allows many devices to run seamlessly on a single shared network.

OPEN Alliance

Broadcom developed BroadR-Reach as a proprietary PHY standard to enable longer distances of copper Ethernet connectivity at 100 Mbps. This PHY uses technologies from the Gigabit Ethernet copper, including multilevel PAM3 signaling and better encoding to reduce the bandwidth required on the cable. It also uses echo cancelers to transmit bidirectional data on a single pair. This standard meets automotive EMI requirements because, at 27 MHz bandwidth, it is lower than the 62.5 MHz bandwidth of 100BASE-T. The OPEN Alliance SIG created an open license standard with sponsorship from major players in the automotive market. The industry realized that more than 100 Mbps is needed for video transmission. Still, more is needed to act as a backbone in the car, especially for ADAS and autonomous driving systems. It pushed for creating a task force in the IEEE 802.3 working group (802.3bp) to define a new standard for 1000 Mbps (1 Gbps) over a single twisted pair. This Gigabit Ethernet PHY is known as 1000BASE-T1. A working group for even faster standards, 802.3ch for 2.5/5/10G, was established in 2019. Other task forces are now looking at >10 Gbps standards, such as 802.3cy for 25 G electrical and 802.3cz for 2.5G/5G/10G/25G/50G optical links.

IEEE 802.3 and Open Alliance standards are closely related and complementary. Typically, IEEE clearly defines what the spec should be like. The Open Alliance takes it over to break it out into a whole test setup and procedure details to define “how” you should make that measurement and testing.

Time synchronization

Some car algorithms require a simultaneous sampling of multiple sensors or referencing the time a measurement was taken. As these measurements occur in different nodes, time synchronization must be accurate between all the nodes in the car down to sub-microsecond resolution. IEEE 802.1AS Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks became the synchronized timing standard. This standard uses a profile for IEEE 1588 v2 and introduces simplified, faster methods for choosing a master clock.

Time-Triggered Ethernet

Some time-sensitive controls require communication latency in the single-microsecond range so that the controller can quickly get the sensor readings or control a time-critical function. In traditional Ethernet, a new packet would have to wait until an existing packet is complete, taking hundreds of microseconds, even at gigabit speeds. The IEEE 802.3br (Interspersing Express Traffic) task force is addressing this problem by developing a system in which a high-priority packet (called an express packet) can interrupt existing packets, and the interrupted packets continue after transmission of the express packet.

AV bridging

ADAS relies on getting data from cameras and other sensors promptly. Unlike watching videos on a computer, where buffering is used to compensate for the unreliable timing of the network, automotive AV systems require both controlled latencies and guaranteed bandwidth. The Time-Sensitive Networking Task Group provides the specifications that allow time-synchronized low-latency streaming services.

Multi-gigabit Optical Automotive Ethernet

Optical Automotive Ethernet is a technology that uses optical fibers to transmit data between different systems in a vehicle. The IEEE P802.3cz task force completed with the approval of IEEE std 802.3cz-2023 in March 2023. The IEEE 802.3cz technology provides the optical connectivity solution for up to 50 Gbps single lane connections with 4 inline connectors and a maximum length of 30 meters in the car over OM3 multimode fiber. The key advantages for optical automotive ethernet include enhanced electromagnetic compatibility (EMI/EMC), very little signal attenuation over a long run, high data ratees, and low cable harness weight.

Energy-efficient Ethernet

Energy-efficient Ethernet is a specific protocol that helps reduce vehicle power consumption by turning off network devices or entire subnets when the engine is off. It includes an energy-optimized Ethernet implementation when the engine is running.

Comprehensive test requirements for successful implementation

Automotive Ethernet engineers deal with ordinary high-frequency board design challenges, including signal noise, signal quality, cross talk, reflection, impedance matching, and DC power integrity.

Automotive Ethernet also requires comprehensive physical layer, protocol, conformance testing, security, and harness testing to ensure successful implementation and reliable operation.

There are three testing points for physical-layer compliance, as described in Figure 3:

- the transmitter with protocol trigger and decode.
- the link segment, including harnesses and connectors.
- the receiver

Note that although this figure shows linear Tx on one end and Rx on the other, automotive Ethernet is a bidirectional bus and both ends of the link will need to be tested for both Tx and Rx.

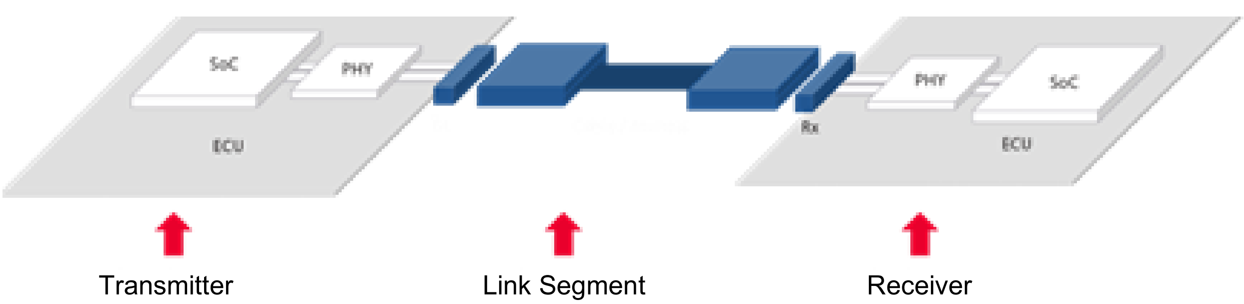

Figure 3. Automotive Ethernet physical-layer test points

Transmitter testing

The transceiver tests are like other high-speed digital PHY characterization solutions. Choosing a testing solution that offers a protocol trigger and decode software package that will view the data traffic and the protocol layer is imperative, saving debugging time early in design. All the compliance tests need to be prepackaged in setup, configuration, and reporting so that designers can focus on their core mission and deadlines.

Receiver testing

A robust receiver compliance solution providing bit error rate, SQI, and Alien crosstalk noise rejection testing should automatically configure all the necessary test equipment, simplifying and speeding up the overall test process:

- Simplify receiver compliance testing.
- Automatically configure all the required equipment to reduce test time.
- Show connections to the device under test in diagrams.
- Create a printable pass / fail HTML report with margin analysis.

Link segment testing

A complete link segment compliance solution needs to support a cable test, connector test, communication channel test, connector test for crosstalk, and cross-talk test across the entire communication channel.

A Better Way Forward

Autonomous driving and ADAS will benefit society but will present engineers with new test challenges. With increasing demands for high data rates, bandwidth, data security, and future readiness, automotive Ethernet offers advanced capabilities. It overcomes the limitations of traditional automotive serial buses for in-vehicle electronic systems connections and communications.

Keysight helps engineers implement automotive Ethernet technologies by providing full test solutions. Following IEEE standards for transmitters, link segments, and receivers, the AE6000 software suite automates otherwise tedious and manual conformance and interoperability testing. These test solutions offer hardware and accessories to ensure you have all the necessary components to get up and running quickly. D9020AUTP oscilloscope software supports customizable, automated triggering on and decoding of automotive Ethernet 10BASE-T1S and 100/1000BASE-T1 protocols to aid in debugging and analyzing protocol-level, time-correlated data. Learn more about each product in our automotive Ethernet solution below:

- AE6000T Automotive Ethernet Tx Compliance Solution
- AE6000L Automotive Ethernet Channel Testing Solution
- AE6000R Automotive Ethernet Rx Compliance Solution
- D9020AUTP Automotive Ethernet Triggering and Decode (10BASE-T1S and 100/1000BASE-T1) for Infiniium Oscilloscopes
- Infiniium MXR B-Series Oscilloscopes

For more information, please visit www.keysight.com/find/autonomous-driving