



AUTOMOTIVE ETHERNET

BEGINNER'S GUIDE

AUTHOR:

Steffen Gugenhan | Senior Product Manager | STAR ELECTRONICS GmbH & Co. KG

STAR COOPERATION GmbH

STAR COOPERATION GmbH is certified as per DIN EN ISO 50001:2015, ISO 9001:2015 and ISO 14001:2015.
Otto-Lilienthal-Straße 5 | 71034 Boeblingen, Germany | Phone +49 (0)7031 6288-300 | www.star-cooperation.com

AUTOMOTIVE ETHERNET - BEGINNER'S GUIDE

Ethernet has been used in vehicles for various functions such as diagnosis and flashing of ECU's. Ethernet is now increasingly being used as a bus system in vehicles to implement other functions, such as driver assistance. Development engineers in the automotive industry are dealing much more with this technology in their day-to-day work. **What is there to know when dealing with vehicle Ethernet bus systems? What practical issues may arise? And what possible solution strategies are available?**

LAYER FOR LAYER - THE FOUNDATION OF AUTOMOTIVE ETHERNET

A basic understanding of its operation is necessary to recognize and avoid pitfalls when dealing with Ethernet. This can be most simply explained by using a layer model. The OSI Layer Model consists of seven layers, each with well-defined tasks and clear interfaces to the next layer. This model ensures that, ideally, individual layers may be exchanged. Alternatively, one can use the DOD Layer Model, which is explained in greater detail here. This groups together several OSI Layer Model layers and reduces complexity.

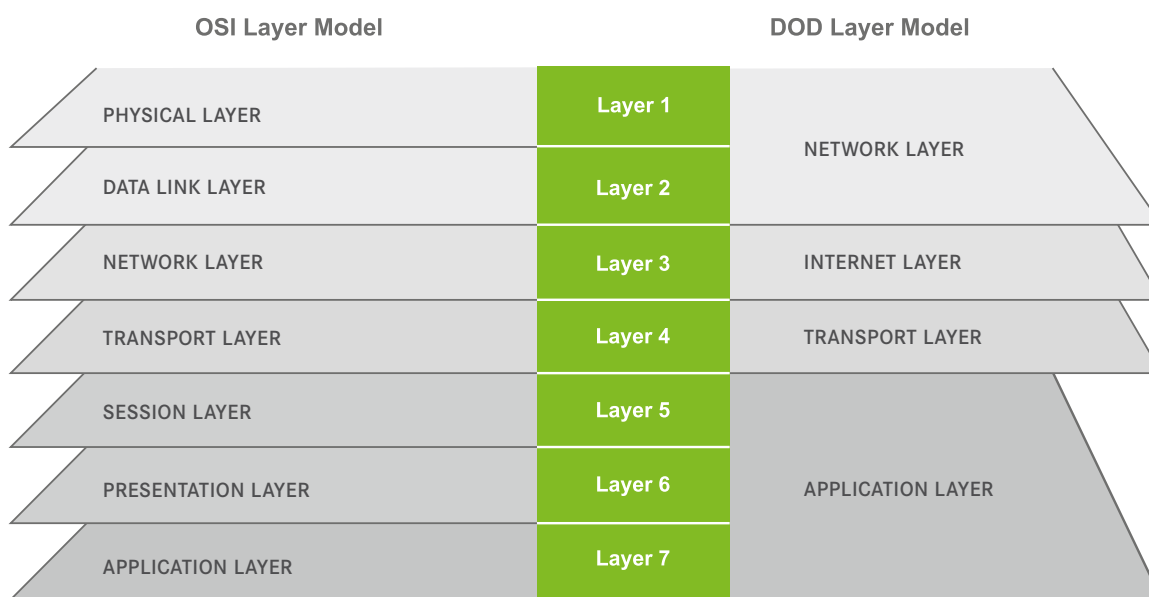


Fig.: OSI vs. DOD Layer Model

LAYER 1: THE NETWORK LAYER

In the DOD layer model **NETWORK LAYER**, data transfer is described on a physical level. This includes the transmission medium, which is the physical layer, as well as the corresponding access protocol, in this case Ethernet.

Special physical layers were developed in accordance with the requirements of the automotive industry for the widespread use of Ethernet in vehicles. One of the considerations was the wish to use unshielded twisted-pair cables for weight and cost reasons and to ensure a high degree of insensitivity to electromagnetic interference.

In the layer model, the physical layer can be exchanged as required without the layers above it being affected or needing to be adjusted. This enormously increases flexibility. The use of a different physical layer makes it possible to easily implement higher data rates within the network. The following physical layers are and will be used in vehicles:

100 Base TX
▪ Standard physical layer (from PC field)
▪ Use in vehicles: Diagnostic interface
▪ 4 wires, shielded cable (STP)
▪ Maximum data transfer rate: 100 Mbps
▪ Maximum cable length: 100 m
▪ Full duplex
100 Base T1
▪ Special physical layer for the automotive sector
▪ Use in vehicles: Networking of ADAS ECU's
▪ 2 wires, unshielded cable (UTP)
▪ Maximum data transfer rate: 100 Mbps
▪ Maximum cable length: 15m
▪ Full duplex
1000 Base T1 (under development)
▪ Special physical layer for the automotive sector
▪ Use in vehicles: Networking of ADAS ECU's and infotainment systems
▪ 2 wires, unshielded (UTP)
▪ Maximum data transfer rate: 1,000 Mbps
▪ Maximum cable length: 15m
▪ Full duplex

Table of the different physical layers in vehicles and their properties

Another important element in the network layer is corresponding network topology. Switches are used to connect Ethernet controllers to one other. These are coupling elements with a certain number of coupling interfaces that actively make a forwarding decision. That is, after a learning phase is completed, the switch only forwards data packets to the coupling interface where the receiver is located. This reduces traffic throughout the network.

In contrast, a hub, which could also be used as a coupling element, generally forwards all traffic to all coupling interfaces. This creates unnecessarily high network traffic volumes. Finally, routers acting as coupling elements provide the ability to transfer traffic from one network segment to another, e.g. from one VLAN to another VLAN. In contrast, a switch only forwards data within its own network segment.

In most cases, star topology is used as network topology. This can be extended by additional switches.

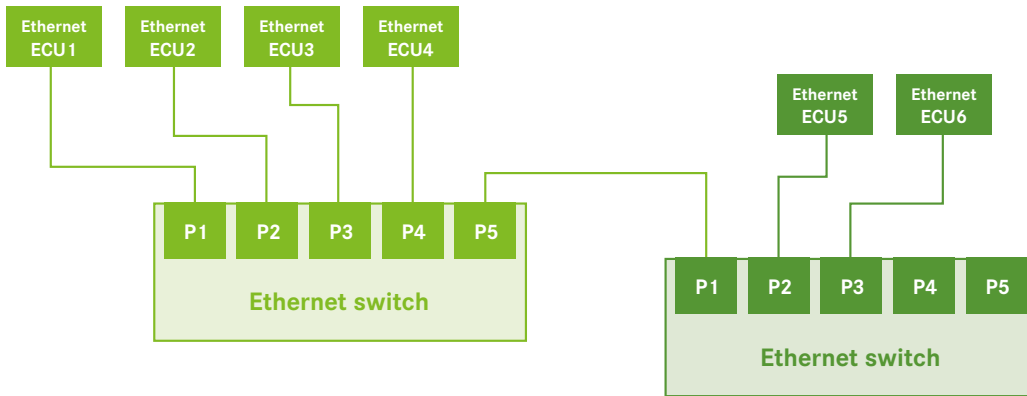


Fig.: Star topology

Data can be sent and received once a physical layer has been selected and a network topology has been established with appropriate switches. The preferred option here is Ethernet, a packet-switched network where data is divided into packets and provided with a frame. According to 802.3 tagged standard, an Ethernet packet framework is subject to the following structure:

Preamble	SFD	Ethernet frame: min. 68 bytes/max. 1522 bytes					
101010..	10101011	Destination MAC address	Source MAC address	VLAN Tag	Type	Data	CRC check sum
8 Bytes		6 Bytes	6 Bytes	4 Bytes	2 Bytes	42-1497 Bytes	4 Bytes

Table: Ethernet packet framework structure

Each Ethernet node is assigned a 6-byte hardware address which identifies the corresponding network node. The destination or source address defines the sender and receiver of the packet. In order to perform a broadcast, which means to send a packet that is then received by a network node, a packet is sent to the destination MAC address 0xFF-FF-FF-FF-FF.

VLAN TAG

A physical network can be subdivided into additional logical subnetworks using the VLAN tag.

Type

The “Type” field defines the type of protocol used in the next layer. If IP were added as the next protocol, the type would be 0x0800 for IPv4 or 0x866D for IPv6.

Data

The “Data” field contains the data of the next layer(s). Each subsequent layer inserts its data into the designated data packet field. Thus, an Ethernet frame forms the basis or lowest level for any data that is transferred over an Ethernet network. In the following layer, the Internet layer, the IP protocol is inserted into the data field of the Ethernet frame.

LAYER 2: THE INTERNET LAYER

Internet layer protocols are tasked with ensuring network-wide addressing of the network participants, regardless of the transmission medium. IP is used for this, making it the first layer independent of the transmission medium. There are two different versions of IP: IPv4 and IPv6, which differ by the size of the address space. IPv4 is explained in more detail here.

The IP data packet is inserted in the “Data” area of the Ethernet frame. Besides the header, it contains a “Data” field for the data in the next layers. The IP header contains, among other elements, the destination IP address and the source IP address. The destination MAC address is required in addition to the destination IP address (see layer 1) in order to be able to send an Ethernet frame. If this is not (yet) known to the sender, it can send an ARP request with the broadcast MAC address 0xFF-FF-FF-FF and will then receive an ARPReply with its MAC address from the corresponding network node.

Preamble	SFD	Ethernet frame: min. 68 bytes/max. 1522 bytes					
101010..	10101011	Destination MAC address	Source MAC address	VLAN Tag	Type	Data	CRC check sum
8 Bytes		6 Bytes	6 Bytes	4 Bytes	2 Bytes	42- 1497 Bytes	4 Bytes

Data		
IP header		IP data
Source IP address	Destination IP address	Data

Table: IP embedding in the Ethernet frame

LAYER 3: THE TRANSPORT LAYER

Transport protocols are contained in the transport layer. These ensure that the Ethernet frame data reaches the corresponding process on the Ethernet node. Both TCP and UDP use so-called “ports” for this purpose. Each transmission therefore contains source and destination ports.

UDP is a connectionless transport protocol where there is no flow control. This means that the package will be sent without a confirmation from the recipient that the package has arrived. In contrast, TCP is a connection-oriented transport protocol in which a connection is set up and every message that is received is also acknowledged. This way, the sender can be sure that the message has been received by the receiver.

Preamble	SFD	Ethernet frame: min. 68 bytes/max. 1522 bytes					
101010..	10101011	Destination MAC address	Source MAC address	VLAN Tag	Type	Data	CRC check sum
8 Bytes		6 Bytes	6 Bytes	4 Bytes	2 Bytes	42- 1497 Bytes	4 Bytes

Data		
IP header		IP data
Source IP address	Destination IP address	Data

UDP or TCP data		
Source port	Destination port	Data

Table: TCP or UDP embedding in an Ethernet frame

LAYER 4: THE APPLICATION LAYER

Actual application protocols are contained in this final layer. They are located in the “Data” field of the transport protocol. This is also the layer where the vehicle data for the individual ECU's are transmitted. A protocol such as Some/IP or a proprietary protocol is used for the networking of ECU's in vehicles.

When all layers have been passed through, the Ethernet frame is complete and ready to be sent. The sender starts with layer 4 (application layer) and the message is passed from layer 4 to layer 1. Each layer adds its information until the message is finally sent by the physical layer. On the receiver side, the process runs in opposite directions, meaning from layer 1 to layer 4. Each layer removes information, leaving only the application protocol payload in the end.

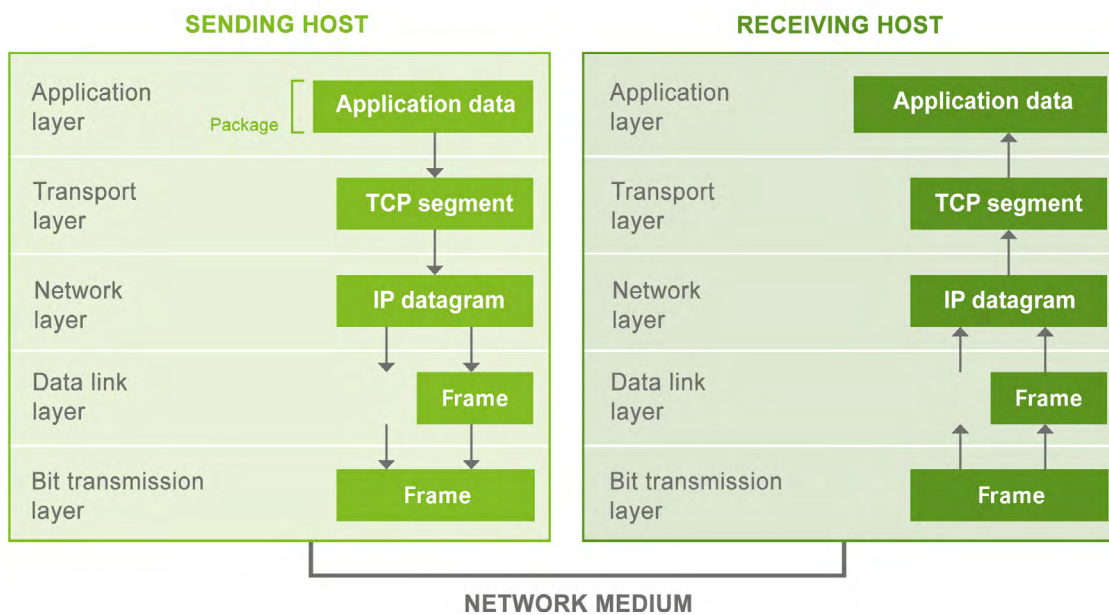


Fig.: Layer model for data transmission

Even though the complexity of networking has greatly increased due to the use of Ethernet as vehicle bus system, there are also decisive advantages. For example, data transmission rates can be flexibly increased by using a different physical layer without having to change the underlying technology. Therefore, development and test engineers should be prepared to use this technology in the long term.

USE CASES AND SOLUTION STRATEGIES | TAP'S, SWITCHES OR BOTH?

The foundations have now been laid for the use of Ethernet in vehicles. But what are concrete use cases and solution strategies?

MEASURING ETHERNET TRAFFIC

In previous vehicle bus systems, such as CAN or FlexRay, traffic could be tapped at any point by means of a Y-cable or a star and all the data traffic of the bus system be read. It is not that easy in an Ethernet network. Individual nodes are connected to each other via switches, so the only traffic that can be tapped at a star arm, is the traffic that is also needed at this point. In addition, the signal quality deteriorates in case of a direct signal pick-up using a Y-cable and thus the connection breaks up. This is especially true for 100 Base T1 and 1000 Base T1 connections. So for automotive-specific physical layers, measurements using Y-cables are not an option.

But there are other ways to tap Ethernet bus traffic. Each has its own respective advantages and disadvantages.

One way is to use a TAP (Test Access Port). This piece of hardware establishes a passive access point to a network connection. An Ethernet line (point-to-point connection) is disconnected and the TAP placed in between. Traffic is routed through one or two separate standard Ethernet interfaces. A TAP is the only way to guarantee tapping of the full traffic in this point-to-point connection, since it has no influence on the traffic itself. However, a TAP causes a certain amount of propagation delay, which could affect the overall system. TAP's for optical data connections are not affected by this disadvantage, but these are not yet used in vehicles. These propagation delays add up when a point-to-point connection is tapped using a TAP in a network topology where multiple switches are connected in series.




Fig.: Data tapping of an Ethernet ECU4 by means of a TAP

Advantages of measuring with a TAP:

- Ensures that all data transmitted via this connection is tapped.
- Easy configuration using Plug & Play.

Disadvantages of measuring with a TAP:

- Additional hardware is required.
- Slight propagation delay of the data as it passes through the TAP. This could have an impact on the overall system, especially if your system is using time synchronization mechanisms, such as 802.1AS.
- Relatively high wiring requirements. If multiple point-to-point connections need to be tapped, things can quickly become quite cluttered.

Another option is to measure at an existing switch. There are already one or more switches in each network topology that make up the entire structure. If there are free ports on the existing switches, they can be used to tap traffic. The free port can be configured as a mirror port or SPAN port depending on the switch type. One must specify in the switch configuration which switch ports should be mirrored to the free port.




Fig.: Tapping Ethernet traffic using a mirror port

Benefits of measuring with an existing switch:

- No additional hardware required.
- No additional latency or runtime modifications of the frames at low load.
- Minimal wiring requirements.

Disadvantages of measuring with an existing switch:

- Packages that are subjected to CRC errors at OSI layer level 2 are discarded by the switch and thus not mirrored on the mirror port.
- The use of mirror ports adds loads to the switch processor, which can cause data loss on the mirror port in case of busy networks.
- Changes in timing behavior with frame interaction are possible.
- Configuring the switch can be complex and expensive, depending on the type of switch.

If no ports or TAP are available on the existing switches, an additional switch for measuring data traffic can be considered as yet another option. As with a TAP, the line is disconnected and the switch placed in between. Data traffic is then mirrored through the mirror port to be configured. This method is subject to the most disadvantages, however, since the respective disadvantages of using an existing switch and a TAP are combined. This is because an additional switch also generates a propagation delay and thus modifies the runtime behavior of the system. Likewise, as with a TAP, additional wiring is needed.




Fig.: Measuring Ethernet ECU5 traffic with an additional switch

And finally, there is the use case where Ethernet traffic is measured in parallel to other bus systems, such as CAN or FlexRay. In most cases, a uniform, synchronized timestamp must be generated for the measured bus systems' data packets.

If the data traffic is routed from the mirror port of the switch or the outputs of the TAP to a standard network interface card of the PC, one must consider that a relatively inaccurate timestamp is generated in Windows. For example, when using a separate card for recording the CAN traffic and the default network interface card of the PC for recording the Ethernet traffic, the generated timestamps are not synchronized and the Ethernet timestamps will be relatively inaccurate. In this case, it may be necessary to choose a bus interface card that has both a CAN and an Ethernet interface that generates an accurate and synchronized timestamp for both interfaces.




Fig.: Measuring Ethernet and CAN traffic with an exact timestamp

There is no such thing as the perfect method for tapping traffic on an Ethernet network. Every one of the options presented has certain advantages and disadvantages. Therefore, you should first think about what data you require for your analysis and decide on that basis which method is the most appropriate to capture that data.

COMMISSIONING AN ETHERNET ECU IN THE LABORATORY

In some cases, it might be advisable to run a variety of tests in the lab without having the entire network topology available. A residual bus simulation is necessary if an Ethernet ECU is commissioned for this purpose. This simulates and generates the data traffic that this ECU would need or receive in a vehicle and sets the ECU to a regular operating state. This allows the ECU to be set to a regular operating state. Both PC-based and autonomous embedded systems may be used to provide residual bus simulation.

Software and a corresponding interface with the appropriate physical layer are necessary to connect the ECU to the PC in the case of PC-based systems. To achieve this, one might use a USB or PCIe connection or converters that convert the data traffic from the standard PC Ethernet card to the corresponding physical layer. The residual bus simulation is performed on the PC and the data is output to the ECU via the interface.

One does not require a PC for the actual operation of the residual bus simulation if an autonomous embedded system is used. The PC is then merely used to configure the residual bus simulation. The ECU to be tested is connected to the embedded system on which the residual bus simulation is performed. Since most PC-based systems use Windows as an operating system and Windows is not real-time capable in its standard configuration, real-time behavior of the embedded solutions is significantly better in these cases. This may be relevant for the tests to be performed, depending upon the application. Most manufacturers provide an easy-to-use user interface for both PC-based systems and autonomous embedded systems. This makes it very easy to configure the residual bus simulation and to put it into operation. Manufacturers also offer solutions for manipulating runtime data, such as GUI's or interfaces.

INTEGRATION OF AN ETHERNET ECU INTO AN OLDER MODEL VEHICLE

It is quite common to integrate an Ethernet ECU of a future vehicle model into a vehicle of an existing model line for testing purposes. The new Ethernet ECU must be connected to a FlexRay or CAN bus or to an Ethernet network with an incompatible communication matrix. Both cases are possible. A gateway may be used for both that performs two tasks: On the one hand, it simulates the components or signals that are not actually present on the Ethernet ECU or the vehicle, so that both the vehicle and the Ethernet ECU can be set into a regular operating state. On the other hand, it routes the required signal values between the two components so that the vehicle can completely operate with the foreign ECU.

The gateway's start-up time is relevant here. After powering on, the required signals are available on the bus in time. Equally important is the ability to support the vehicle network management. If both features are supported, the gateway solution can be solidly integrated into the vehicle and the gateway boots up automatically when the vehicle is started.



Fig.: FD-L as gateway between vehicle and ECU

CONCLUSION AND OUTLOOK

Tool manufacturers are quickly adapting to this new environment since Ethernet will be the vehicle bus system of the future, even if some pitfalls and peculiarities for targeted tests and testing procedures must be considered when dealing with the application of Ethernet in vehicles. Of course, there will always be features that still need to be implemented, such as new database format support. In principle, all the tools required to handle vehicle Ethernet are already available today.

AUTHOR:

Steffen Gugenhan | Senior Product Manager | STAR ELECTRONICS GmbH & Co. KG

CONTACT:

Dominik Dussling | Manager STAR MARKETING | Phone +49 7031 6288-207
dominik.dussling@star-cooperation.com