

EtherCAT – The Ethernet Fieldbus



EtherCAT[®] 
Technology Group



3 A bus system might not be everything ...

4 EtherCAT at a glance

6 EtherCAT Technology Group

8 Why use EtherCAT?

10 The technology in detail

10 EtherCAT: based on Ethernet technology

11 How does EtherCAT work?

12 The EtherCAT protocol

14 Flexible topology

16 EtherCAT P: communication and power in one cable

18 Distributed Clocks for high-precision synchronization

20 Diagnostics and error localization

21 High availability requirements

22 System overview

24 Safety over EtherCAT

26 Communication profiles

29 Plant-wide communication with the EtherCAT Automation Protocol (EAP)

31 Integration of other bus systems

32 Powering the digital transformation with EtherCAT, Industrie 4.0 and IoT

34 Implementing EtherCAT interfaces

36 Implementing a master device

38 Implementing a slave device

40 Conformance and certification

43 Contact

A bus system might not be everything ...

... but the machine is nothing without it!

Not only is it a central part of the system architecture, but its performance also determines whether the entire system is able to reach its full performance. The bus system is also a key factor in determining system costs, commissioning time, and robustness. That's why a good engineer selects the right bus technology as the first step to system design.

We've written this brochure to introduce you to EtherCAT, the Ethernet fieldbus. Not only will you get to know EtherCAT – you will also learn what makes EtherCAT the fastest Industrial Ethernet

standard. This brochure also introduces the EtherCAT Technology Group (ETG), the world's largest fieldbus organization. Most importantly, we hope to convey why EtherCAT is the right choice for your application. If you have any questions, feel free to contact us. We're passionate about EtherCAT and look forward to hearing from you.

On behalf of the
EtherCAT Technology Group team,
Martin Rostan, Executive Director,
EtherCAT Technology Group



Martin Rostan, Executive Director,
EtherCAT Technology Group



The worldwide ETG team
at a global strategy meeting.

EtherCAT is a high-performance, low-cost, easy to use Industrial Ethernet technology with a flexible topology. It was introduced in 2003 and has been an international standard since 2007. The EtherCAT Technology Group promotes EtherCAT and is responsible for its continued development. EtherCAT is also an open technology: anyone is allowed to implement or use it.

How it works

EtherCAT's key functional principle lies in how its nodes process Ethernet frames: each node reads the data addressed to it and writes its data back to the frame all while the frame is moving downstream. This leads to improved bandwidth utilization (one frame per cycle is often sufficient for communication) while also eliminating the need for switches or hubs.

Network performance

The unique way EtherCAT process frames makes it the fastest Industrial Ethernet technology; no other technology can top EtherCAT's bandwidth utilization or the corresponding performance.

Flexible topology

In addition to its speed, an EtherCAT network can support up to 65,535 devices without placing restrictions on their topology: line, bus, tree, star – or any combination thereof. Fast Ethernet physics allows two devices to be up to 100m (330 ft.) apart, and greater distances are possible with the use of fiber optics. EtherCAT also has additional features that offer further topological flexibility, such as Hot Connect and Hot Swap for devices, and added redundancy through a ring topology.

It's versatile

EtherCAT is suitable for both centralized and decentralized system architectures. It can support master-to-master, master-to-slave and slave-to-slave communication as well as incorporate subordinate fieldbuses. At the factory-level, the EtherCAT Automation Protocol has communication covered – all with the existing infrastructure.

It's easy

When compared to a classic fieldbus system, EtherCAT is the obvious choice: node addresses can be set automatically, there's no need for network tuning, and onboard diagnostics with fault localization make pinpointing errors a snap. Despite these advanced features, EtherCAT is also easier to use than Industrial Ethernet: there are no switches to configure, and no complicated handling of MAC or IP addresses is required.

It's low-cost

EtherCAT delivers all of the advantages of Industrial Ethernet at fieldbus prices. How? Firstly, EtherCAT doesn't require any active infrastructure components. The master device doesn't require a special interface card and the slave devices use highly-integrated, cost-effective chips available from a variety of suppliers. Additionally, there's no need for costly IT experts to commission or maintain the system.

Industrial Ethernet

EtherCAT also supports common internet technologies without jeopardizing the network's real-time capability. Its Ethernet over EtherCAT protocol transports FTP, http, TCP/IP, etc.

Functional safety

Safety over EtherCAT is just like EtherCAT itself – lean and fast. Functional safety is built directly into the bus with options for both centralized and decentralized safety logic. Thanks to the Black Channel approach, it is also available for other bus systems.

Open technology

EtherCAT is an internationally standardized open technology, meaning anyone is free to use the technology in a compatible form. However, being an open technology doesn't mean that anyone can arbitrarily change EtherCAT to suit his or her needs. This would be the end of interoperability. The EtherCAT Technology Group, the world's largest fieldbus organization, is responsible for the further development of EtherCAT so it remains both open and interoperable.

It's proven in use

EtherCAT is currently in use across the globe in an unrivaled range of applications. EtherCAT is used in machine control, measurement equipment, medical devices, automobiles and mobile machines, as well as in innumerable embedded systems.



ETG trade show booths present the technology itself as well as the EtherCAT product variety.



ETG holds regional member meetings on a regular basis.

The EtherCAT Technology Group keeps EtherCAT technology open for all potential users. It brings EtherCAT device manufacturers, technology providers, and users together to further the technology. It provides multiple Technical Working Groups where experts carefully work on various specific aspects of EtherCAT. All of these activities are focused on one common goal: keeping EtherCAT stable and interoperable. That's why there is only a single version of EtherCAT, and not a new version each year.

The ETG holds multiple Plug Fests in Europe, Asia, and America each year. The Plug Fests bring EtherCAT device developers together to test and ensure device interoperability. Using the official EtherCAT Conformance Test Tool, each manufacturer conformance tests its EtherCAT devices prior to their release. The ETG awards the manufacturer a Conformance Certificate following a successful test in an accredited test lab.

The ETG also holds international seminars and workshops and represents EtherCAT at tradeshows throughout the globe. It also provides product guides, joint tradeshow booths, and seminar exhibits to help its members market their EtherCAT products.

The ETG has the largest number of members out of any fieldbus organization in the world. The list of members can be found on its homepage. However, the decisive factor is not how many members there are, but how active the members are in the ETG. Both the number and variety of EtherCAT devices is unparalleled, and EtherCAT's adoption rate across Europe, Asia, and America is outstanding.

Become a member

ETG membership is open to all companies, whether device manufacturers or users. ETG members:

- Receive access to technical specifications and the developer's forum
- Contribute to the further development of EtherCAT in the way of Technical Working Groups
- Receive implementation support from their local ETG office
- Receive free software stacks, software tools, and access to additional development products
- Are permitted to use the EtherCAT and ETG logos
- Display their EtherCAT products and services in the official EtherCAT Product Guide, at tradeshows, and at ETG events.

Bylaws, registration form, and additional information are available via info@ethercat.org and www.ethercat.org.

International standardization

The EtherCAT Technology Group is an official partner of the IEC. Both EtherCAT and Safety over EtherCAT are IEC-Standards (IEC 61158 and IEC 61784). These standards not only include the lower protocol layers, but also the application layer and device profiles, e.g. for drives. SEMI™ (Semiconductor Equipment and Materials International) has accepted EtherCAT as a communication standard (E54.20) for the semiconductor industry. The various Task Groups in the ETG Semiconductor Technical Working Group (TWG) have defined industry specific device profiles and implementation guidelines.

The EtherCAT Specification is available in English, Japanese, Korean, and Chinese.

Global activities

The EtherCAT Technology Group is active worldwide. Experts in the ETG offices in Germany, China, Japan, Korea, and the USA support ETG members before, during and after implementation.

The technology is maintained in several Technical Working Groups, which define enhancements and also uniform device behavior by means of device profiles and implementation guidelines. All members are encouraged to contribute actively to the TWGs.



In many countries EtherCAT also is a national standard, such as in Korea and China.

EtherCAT milestones



Why use EtherCAT?

EtherCAT

The unique way that EtherCAT works makes it the clear engineer's choice. Additionally, the following features are particularly advantageous for certain applications.

1. Exceptional performance

EtherCAT is by and large the fastest Industrial Ethernet technology, but it also synchronizes with nanosecond accuracy.

This is a huge benefit for all applications in which the target system is controlled or measured via the bus system. The rapid reaction times work to reduce the wait times during the transitions between process steps, which significantly improves application efficiency. Lastly, the EtherCAT system architecture typically reduces the load on the CPU by 25–30% in comparison to other bus systems (given the same cycle time). When optimally applied, EtherCAT's performance leads to improved accuracy, greater throughput, and thus to lowered costs.

2. Flexible topology

In EtherCAT applications, the machine structure determines the network topology, not the other way around. In conventional Industrial Ethernet systems, there are limitations on how many switches and hubs can be cascaded, which thus limits the overall network topology. Since EtherCAT doesn't need hubs or switches, there are no such limitations. In short, EtherCAT is virtually limitless when it comes to network topology. Line, tree, star topologies and any combinations thereof are possible with a nearly unlimited number of nodes. Thanks to automatic link detection, nodes and network segments can be disconnected

during operation and then reconnected – even somewhere else. Line topology is extended to a ring topology for the sake of cable redundancy. All the master device needs for this redundancy is a second Ethernet port, and the slave devices already supports the cable redundancy. This makes switching out devices during machine operation possible.

3. It's simple and robust

Configuration, diagnostics, and maintenance are all factors that contribute to system costs. The Ethernet fieldbus makes all of these tasks significantly easier: EtherCAT can be set to automatically assign addresses, which eliminates the need for manual configuration. A low bus load and peer-to-peer physics improve electromagnetic noise immunity. The network reliably detects potential disturbances down to their exact location, which drastically reduces the time needed for troubleshooting. During startup, the network compares the planned and actual layouts to detect any discrepancies.

EtherCAT performance also helps during system configuration by eliminating the need for network tuning. Thanks to the large bandwidth, there is capacity to transmit additional TCP/IP together with the control data. However, since EtherCAT itself isn't based on TCP/IP, there is no need to administer MAC-addresses or IP-addresses or to have IT experts configure switches and routers.

4. Integrated safety

Functional safety as an integrated part of the network architecture? Not a problem with FailSafe over EtherCAT (FSoE). FSoE is proven in use through TÜV certified devices that have been on the market since 2005. The protocol fulfills the requirements for SIL 3 systems and is suitable for both centralized and decentralized control systems. Thanks to the Black Channel approach and the particularly lean safety container, FSoE can also be used in other bus systems. This integrated approach and the lean protocol help keep system costs down. Additionally, a non-safety critical controller can also receive and process safety data.

5. Affordability

EtherCAT delivers the features of Industrial Ethernet at a price similar or even below that of a classic fieldbus system. The only hardware required by the master device is an Ethernet port – no expensive interface cards or co-processors are necessary. EtherCAT Slave Controllers are available from various manufacturers in different formats: as an ASIC, based on FPGA, or as an option for standard microprocessor series. Since these inexpensive controllers shoulder all the time-critical tasks, EtherCAT itself doesn't place any performance requirements on the CPU of slave devices, which keeps device costs down. Since EtherCAT doesn't require switches or other active infrastructure components, the costs for these components and their installation, configuration, and maintenance are also eliminated.

For these reasons, EtherCAT is often seen in:

- Robotics
- Machine tools
- Packaging machines
- Printing presses
- Plastic manufacturing equipment
- Presses
- Semiconductor manufacturing machines
- Test benches
- Pick & Place Machines
- Measurement systems
- Power plants
- Substations
- Material handling applications
- Baggage handling systems
- Stage control systems
- Automated assembly systems
- Pulp and paper machines
- Tunnel control systems
- Welding machines
- Cranes and lifts
- Farm machinery
- Offshore applications
- Sawmills
- Window manufacturing equipment
- Building automation systems
- Iron and steel works
- Wind turbines
- Furniture manufacturing equipment
- Milling machines
- Automated guided vehicles
- Entertainment automation
- Medical devices
- Woodworking machines
- Flat glass manufacturing equipment
- Weighing systems
- etc.



The technology in detail

EtherCAT: based on Ethernet technology

EtherCAT is Industrial Ethernet and utilizes standard frames and the physical layer as defined in the Ethernet Standard IEEE 802.3. However, it also addresses the specific demands faced in the automation industry, where:

- There are hard real-time requirements with deterministic response times.
- The system is usually made up of many nodes, each only having a small amount of cyclic process data.
- Hardware costs are even more important than in IT and office applications.

The above requirements make using a standard Ethernet network at the field level practically impossible. If an individual Ethernet telegram is used for each node, the effective data rate sinks significantly for just a few bytes of cyclic process data: the shortest Ethernet telegram is 84 bytes long (including the Inter Frame Gap), of which 46 bytes can be used for process data. For example, if a drive sends 4 bytes of process data for the actual position and status information and receives 4 bytes of data for the target position and control information, the effective data rate for both telegrams sinks to $4/84 = 4.8\%$. Additionally, the drive usually has a reaction time that triggers the transmission of the actual values after receiving the target values. At the end, not much of the 100MBit/s transfer rate remains.

Protocol stacks, such as those used in the IT world for routing (IP) connection (TCP), require additional overhead for each node and create further delays through the stack runtimes.

How does EtherCAT work?

EtherCAT overcomes the difficulties described in the previous section with its high-performing mode of operation, in which a single frame is usually sufficient to send and receive control data to and from all nodes!

The EtherCAT master sends a telegram that passes through each node. Each EtherCAT slave device reads the data addressed to it on the fly, and inserts its data in the frame as the frame is moving downstream. The frame is delayed only by hardware propagation delay times. The last node in a segment or branch detects an open port and sends the message back to the master using Ethernet technology's full duplex feature.

The telegram's maximum effective data rate increases to over 90%, and due to the utilization of the full duplex feature, the theoretical effective data rate is even greater than 100 MBit/s.

The EtherCAT master is the only node within a segment allowed to actively send an EtherCAT frame; all other nodes merely forward frames downstream. This concept prevents unpredictable delays and guarantees real-time capabilities.


The master uses a standard Ethernet Media Access Controller (MAC) without an additional communication processor. This allows a master to be implemented on any hardware platform with an available Ethernet port, regardless of which real-time operating system or application software is used.

EtherCAT slave devices use an EtherCAT Slave Controller (ESC) to process frames on the fly and entirely in hardware, making network performance predictable and independent of the individual slave device implementation.



The EtherCAT protocol

EtherCAT embeds its payload in a standard Ethernet frame. The EtherCAT frame is identified with the Identifier (0x88A4) in the EtherType field. Since the EtherCAT protocol is optimized for short cyclic process data, the use of bulky protocol stacks, such as TCP/IP or UDP/IP, can be eliminated.



EtherCAT in a standard Ethernet frame (according to IEEE 802.3)

To ensure Ethernet IT communication between the nodes, TCP/IP connections can optionally be tunneled through a mailbox channel without impacting real-time data transfer.

During startup, the master device configures and maps the process data on the slave devices. Different amounts of data can be exchanged with each slave, from one bit to a few bytes, or even up to kilobytes of data.

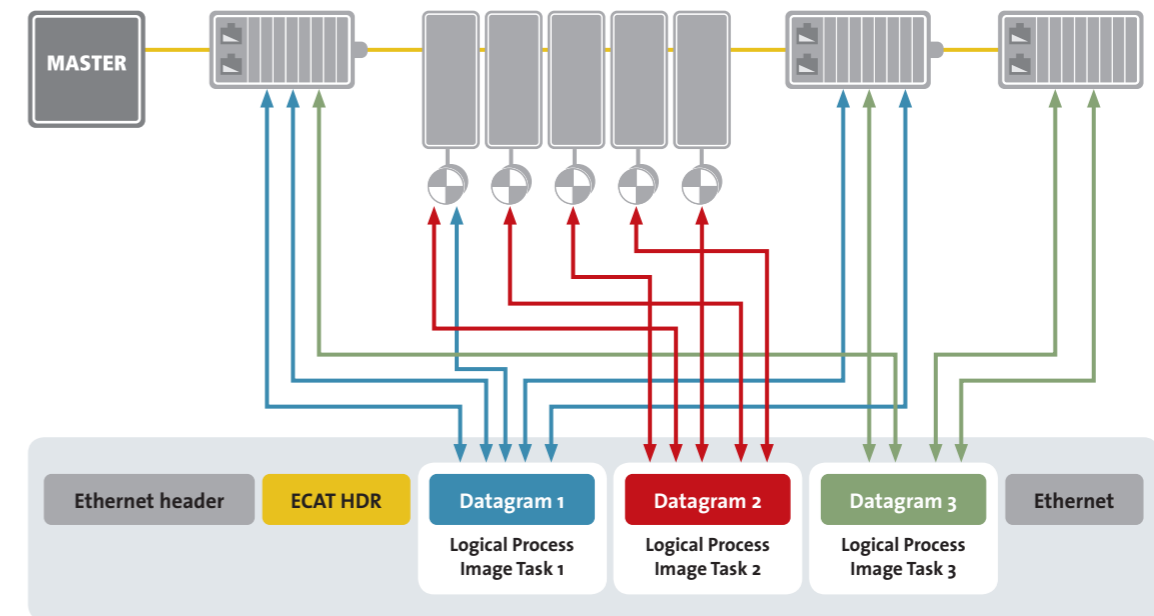
The EtherCAT frame contains the frame header and one or more datagrams. The datagram header indicates what type of access the master device would like to execute:

- Read, write, or read-write
- Access to a specific slave device through direct addressing, or access to multiple slave devices through logical addressing (implicit addressing)

Logical addressing is used for the cyclical exchange of process data. Each datagram addresses a specific part of the process image in the EtherCAT segment, for which 4 GBytes of address space is available. During network startup, each slave device is assigned one or more addresses in this global address space. If multiple slave devices are assigned addresses in the same area, they can all be addressed with a single datagram. Since the datagrams completely contain all the data access related information, the master device can decide when and which data to access. For example, the master device can use short cycle times to refresh data on the drives, while using a longer cycle time to sample the I/O; a fixed process data structure is not necessary. This also relieves the master device in comparison to in conventional fieldbus systems, in which the data from each node had to be read individually, sorted with the help of the process controller, and copied into memory. With EtherCAT, the master device only needs to fill a single EtherCAT frame with new out-

put data, and send the frame via automatic Direct Memory Access (DMA) to the MAC controller. When a frame with new input data is received via the MAC controller, the master device can copy the frame again via DMA into the computer's memory – all without the CPU having to actively copy any data.

In addition to cyclical data, further datagrams can be used for asynchronous or event driven communication.



Inserting process data on the fly


In addition to logical addressing, the master device can also address a slave device via its position in the network. This method is used during network boot up to determine the network topology and compare it to the planned topology.

After checking the network configuration, the master device can assign each node a configured node address and communicate with the node via this fixed address. This enables targeted access to devices, even when the network topology is changed during operation, for example with Hot Connect Groups. There are two approaches for slave-to-slave communication. A slave device can send data directly to another slave device that is connected further downstream in the network. Since EtherCAT frames can only be processed going forward, this type of direct communication depends on the network's topology, and

is particularly suitable for slave-to-slave communication in a constant machine design (e.g. in printing or packaging machines). In contrast, freely configurable slave-to-slave communication runs through the master device, and requires two bus cycles (not necessarily two control cycles). Thanks to EtherCAT's excellent performance, this type of slave-to-slave communication is still faster than with other communication technologies.

Flexible topology

Line, tree, star, or daisy-chain: EtherCAT supports almost all of topologies. EtherCAT makes a pure bus or line topology with hundreds of nodes possible without the limitations that normally arise from cascading switches or hubs.



Flexible topology: line, tree, star, or daisy-chain

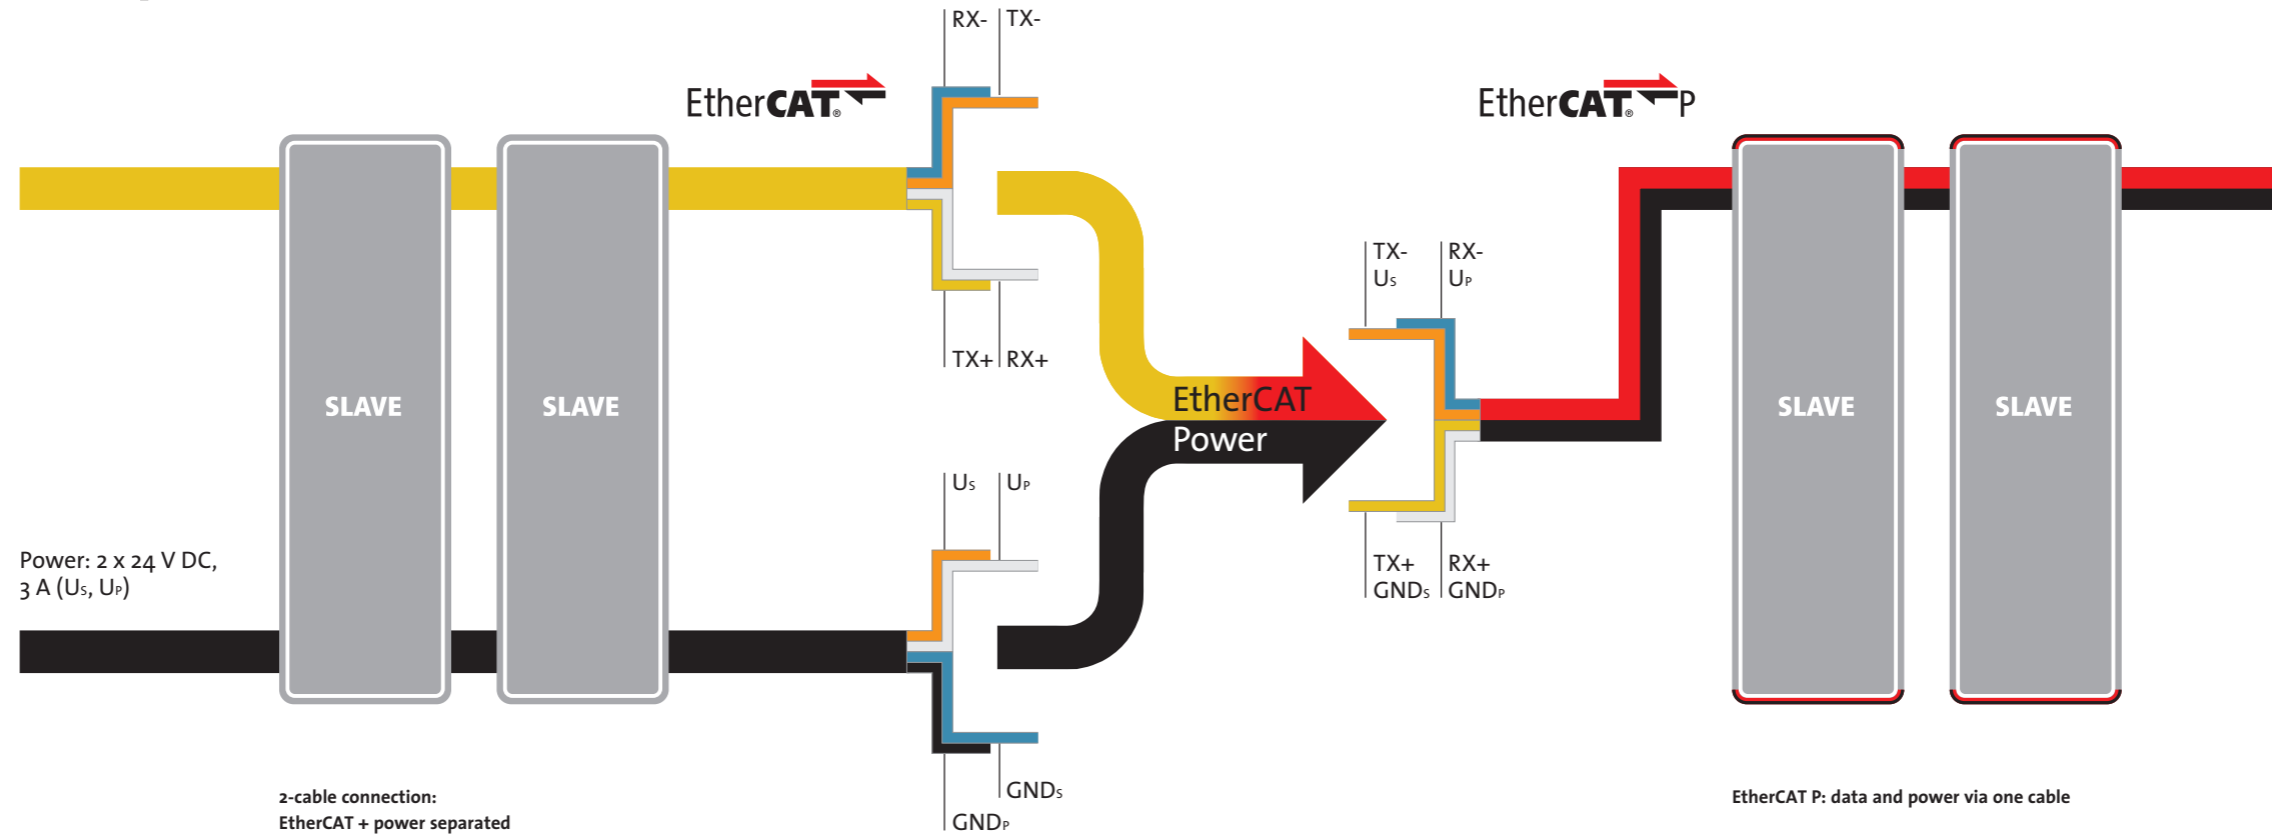
When wiring the system, the combination of lines with branches or drop lines is particularly beneficial: the ports necessary to create branches are directly integrated in many I/O modules, so no additional switches or active infrastructure components are required. The star topology, the Ethernet classic, can also naturally be utilized.

Modular machines or tool changers require network segments or individual nodes to be connected and disconnected during operation. EtherCAT Slave Controllers already include the basis for this Hot Connect feature. If a neighboring station is removed, then the port is automatically closed so the rest of the network can continue to operate without interference. Very short detection times faster than 15 microseconds guarantee a smooth changeover.

EtherCAT offers a lot of flexibility regarding cable types, so each segment can use the exact type of cable that best meets its needs. Inexpensive industrial Ethernet cable can be used between two nodes up to 100m apart in 100BASE-TX mode. Furthermore, the protocol addition EtherCAT P enables the transmission of data and power via one cable. This option enables the connection of devices such as sensors with a single line. Optical fibres (such as 100BASE-FX) can also be used, for example for a node distance greater than 100 m. The complete range of Ethernet wiring is also available for EtherCAT.

Up to 65,535 devices can be connected to one EtherCAT segment, so network expansion is virtually unlimited. Because of the practically unlimited number of nodes, modular devices such as sliced I/O stations can be designed in such a way that each module is an EtherCAT node of its own. Hence, the local extension bus is eliminated; the high performance of EtherCAT reaches each module directly and without any delays, since there is no gateway in the bus coupler or head station any more.

EtherCAT P: communication and power in one cable



EtherCAT P (P = power) is an addition to the previously described EtherCAT protocol standard. It enables not only the transmission of communication data, but also the peripheral voltage via a single, standard four-wire Ethernet cable.

EtherCAT and EtherCAT P are identical in terms of the protocol technology, as the addition exclusively affects the physical layer. No new EtherCAT Slave Controllers are necessary when using EtherCAT P. One could say that EtherCAT P has the same communication advantages as EtherCAT, but also provides the power supply via the communication cable, offering attractive benefits and enhancements for numerous applications.

Two electrically isolated, individually switchable 24 V supplies power the new EtherCAT P devices, available with U_s serving the system and sensors and U_p serving the periphery and actuators. Both voltages, U_s and U_p , are directly coupled into the 100 Mbit/s EtherCAT communication line. Thanks to this power transmission, the user can cascade several EtherCAT P devices and therefore only needs one cable. This facilitates reduced cabling,

more compact, cost-effective wiring, lower system costs and a smaller footprint for devices, equipment and machines.

EtherCAT P is especially interesting for those parts of a machine that are self-contained and often a bit isolated, as they can be supplied with data and power through a single stub cable. Sensors of all types are perfectly suitable for EtherCAT P: a single compact M8 connector enables efficient integration of these field devices into the high-speed network and connects them to the supply voltage. Potential error sources while connecting devices are avoided, thanks to mechanical coding of the connector.

EtherCAT P can be used in the same network together with standard EtherCAT technology. Appropriate rectifier units transform common EtherCAT physics to EtherCAT P by consistently maintaining the Ethernet data encoding. In the same way, a device itself can be supplied with EtherCAT P but, in turn, can also transmit the EtherCAT protocol.


Further information about EtherCAT P can be found here: www.ethercat.org/ethercat-p



Distributed Clocks for high-precision synchronization

In applications with spatially distributed processes requiring simultaneous actions, exact synchronization is particularly important. For example, this is the case for applications in which multiple servo axes execute coordinated movements.


In contrast to completely synchronous communication, whose quality suffers immediately from communication errors, distributed synchronized clocks have a high degree of tolerance for jitter in the communication system. Therefore, the EtherCAT solution for synchronizing nodes is based on such Distributed Clocks (DC).



Completely hardware-based synchronization with compensation for propagation delays.

The calibration of the clocks in the nodes is completely hardware-based. The time from the first DC slave device is cyclically distributed to all other devices in the system. With this mechanism, the slave device clocks can be precisely adjusted to this reference clock. The resulting jitter in the system is significantly less than $1\mu\text{s}$.


Since the time sent from the reference clock arrives at the slave devices slightly delayed, this propagation delay must be measured and compensated for each slave device in order to ensure synchronicity and simultaneousness. This delay is measured during network startup or, if desired, even continuously during operation, ensuring that the clocks are simultaneous to within much less than $1\mu\text{s}$ of each other.



Synchronicity and simultaneousness – scope view of two distributed devices with 300 nodes and 120 m cable length.

If all nodes have the same time information, they can set their output signals simultaneously and affix their input signals with a highly precise timestamp. In motion control applications, cycle accuracy is also important in addition to synchronicity and simultaneousness. In such applications, velocity is typically derived from the measured position, so it is critical that the position measurements are taken precisely equidistantly (i.e. in exact cycles). Even very small inaccuracies in the position measurement timing can translate to larger inaccuracies in the calculated velocity, especially relative to short cycle times. With EtherCAT, the position measurements are triggered by the precise local clock and not the bus system, leading to much greater accuracy.

Additionally, the use of Distributed Clocks also unburdens the master device; since actions such as position measurement are triggered by the local clock instead of when the frame is received, the master device doesn't have such strict requirements for sending frames. This allows the master stack to be implemented in software on standard Ethernet hardware. Even jitter in the range of several microseconds does not diminish the accuracy of the Distributed Clocks! Since the accuracy of the clock does not depend on when it's set, the frame's absolute transmission time becomes irrelevant. The EtherCAT master need only to ensure that the EtherCAT telegram is sent early enough, before the DC signal in the slave devices triggers the output.



Diagnostics and error localization

Experience with conventional fieldbus systems has shown that diagnostic characteristics play a major role in determining a machine's availability and commissioning time. In addition to error detection, error localization is important during troubleshooting. EtherCAT features the possibility to scan and compare the actual network topology with the planned topology during boot up. EtherCAT also has many additional diagnostic capabilities inherent to its system.

The EtherCAT Slave Controller in each node checks the moving frame for errors with a checksum. Information is provided to the slave application only if the frame has been received correctly. If there is a bit error, the error counter is incremented and the subsequent nodes are informed that the frame contains an error. The master device will also detect that the frame is faulty and discard its information. The master device is able to detect where the fault originally occurred in the system by analyzing the nodes' error counters. This is an enormous advantage in comparison to conventional fieldbus systems, in which an error is propagated along the entire party line, making it impossible to localize the source of the error. EtherCAT can detect and localize occasional disturbances before the issue impacts the machine's operation.

Thanks to EtherCAT's unique principle of bandwidth utilization, which is orders of magnitude better than industrial Ethernet technologies that use single frames, the likelihood of bit errors induced by disturbances within an EtherCAT frame is substantially lower – if the same cycle time is used. And, if in typical EtherCAT fashion much shorter cycle times are used, the time required for error recovery is significantly reduced. Thus, it is also much simpler to master such issues within the application.


Within the frames, the Working Counter (WC) enables the information in each datagram to be monitored for consistency. Every node that is addressed by the datagram and whose memory is accessible increments the Working Counter automatically. The master is then able to cyclically confirm if all nodes are working with consistent data. If the Working Counter has a different value than it should, the master does not forward this datagram's data to the control application. The master device is then able to automatically detect the reason for the unexpected behavior with help from status and error information from the nodes as well as the link status.

Since EtherCAT utilizes standard Ethernet frames, Ethernet network traffic can be recorded with the help of free Ethernet software tools. For example, the well-known Wireshark software comes with a protocol interpreter for EtherCAT, so that protocol-specific information, such as the Working Counter, commandos, etc. are shown in plain text.

High availability requirements

For machines or equipment with very high availability requirements, a cable break or a node malfunctioning should not mean that a network segment is no longer accessible, or that the entire network fails.

EtherCAT enables cable redundancy with simple measures. By connecting a cable from the last node to an additional Ethernet port in the master device, a line topology is extended into a ring topology. A redundancy case, such as a cable break or a node malfunction, is detected by a software add-on in the master stack. That's all there is to it! The nodes themselves don't need to be modified, and don't even know that they're being operated in a redundant network.



Inexpensive cable redundancy with standard EtherCAT slave devices

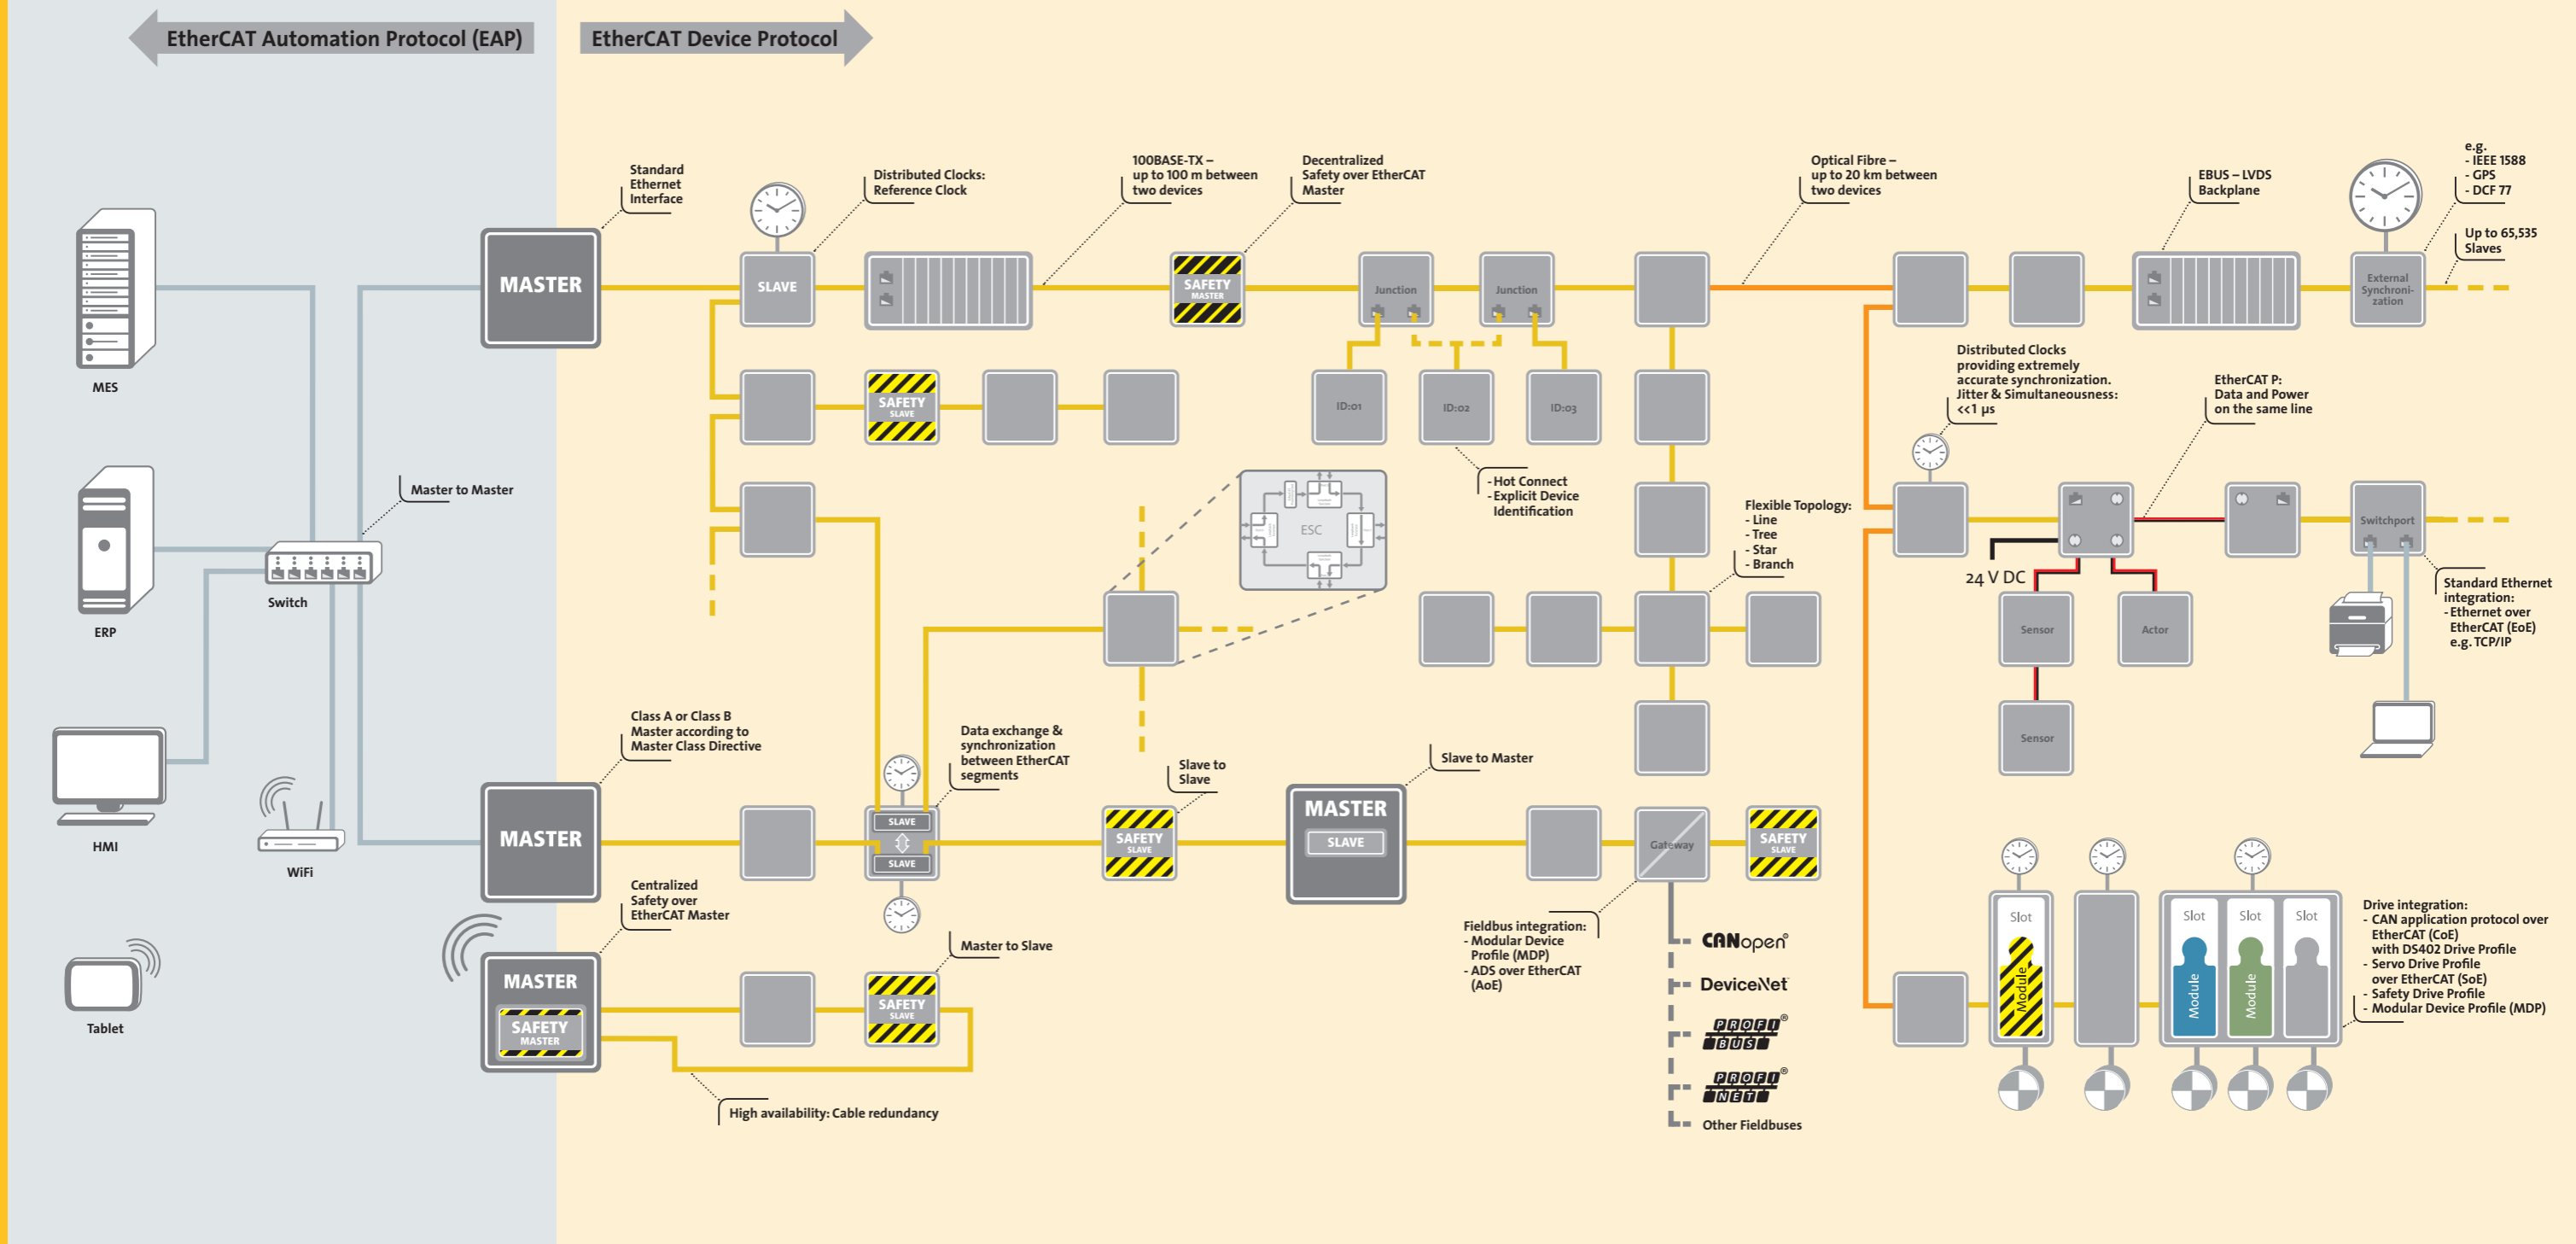
Link detection in the slave devices automatically detect and resolve redundancy cases with a recovery time is less than 15 microseconds, so at maximum, a single communication cycle is disrupted. This means that even motion applications with very short cycle times can continue working smoothly when a cable breaks.

With EtherCAT, it's also possible to realize master device redundancy with Hot Standby. Vulnerable network components, such as those connected with a drag chain, can be wired with a drop line, so that even when a cable breaks, the rest of the machine keeps running.

System overview

EtherCAT plant network

EtherCAT machine control network

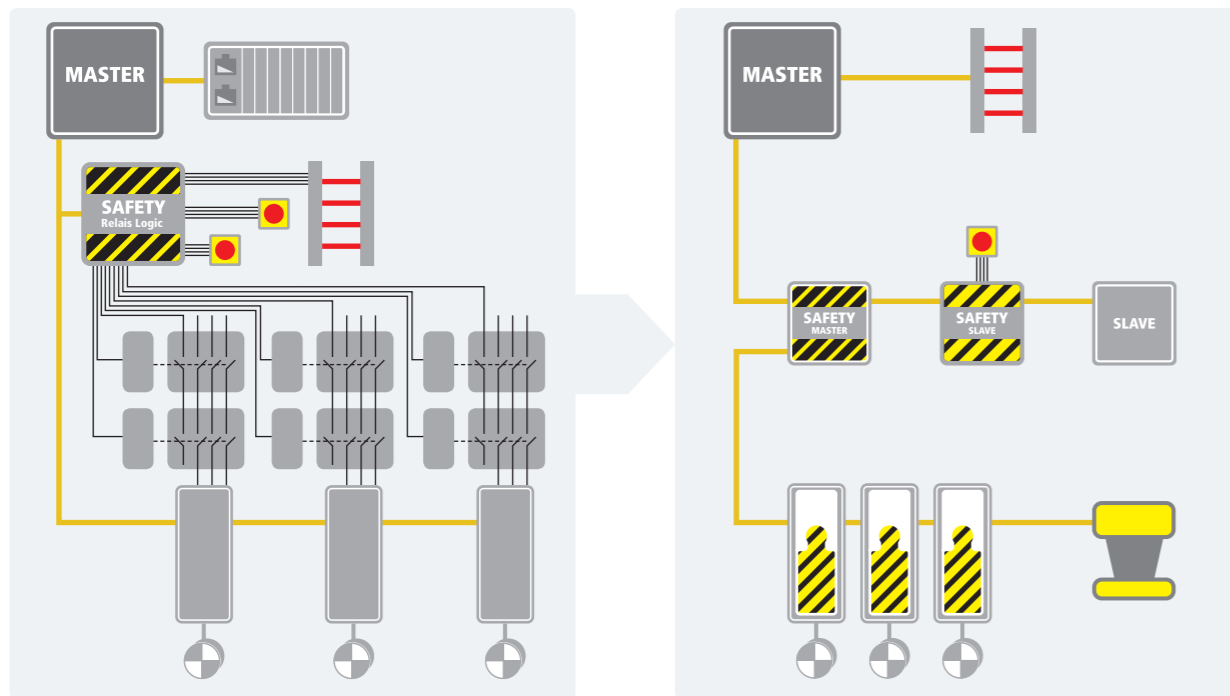


Safety over EtherCAT

Safety over
EtherCAT®

Modern communication systems not only realize the deterministic transfer of control data, they also enable the transfer of safety-critical control data through the same medium. EtherCAT utilizes the protocol Safety over EtherCAT (FSoE = FailSafe over EtherCAT) for this very purpose and so allows:

- A single communication system for both control and safety data
- The ability to flexibly modify and expand the safety system architecture
- Pre-certified solutions to simplify safety applications
- Powerful diagnostic capabilities for safety functions
- Seamless integration of the safety design in the machine design
- The ability to use the same development tools for both standard and safety applications

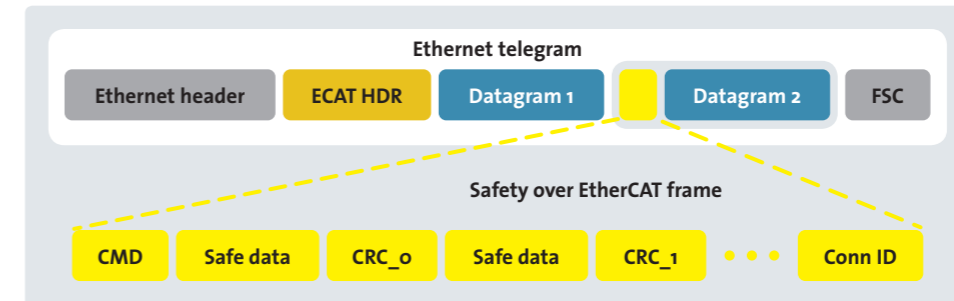


Safety over EtherCAT enables simpler and more flexible architectures than with relay logic.

The EtherCAT safety technology was developed according to IEC 61508, is TÜV certified, and is standardized in IEC 61784-3. The protocol is suitable for safety applications with a safety integrity level up to SIL 3.

With Safety over EtherCAT, the communication system is part of a Black Channel, which is not considered to be safety relevant. The standard communication system EtherCAT makes use of a single channel to transfer both standard and safety-critical data. Safety over

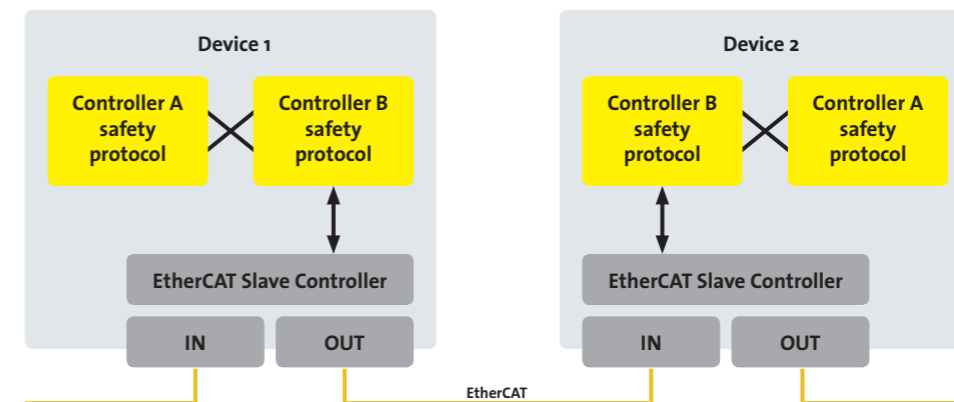
EtherCAT frames, known as safety containers, contain safety-critical process data and additional information used to secure this data. The safety containers are transported as part of the communication's process data. Whether data transfer is safe does not depend on the underlying communication technology, and isn't restricted to EtherCAT; safety containers can travel through fieldbus systems, Ethernet or similar technologies, and can make use of copper cables, fiber optics, and even wireless connections.



The Safety over EtherCAT frame (safety container) is embedded in the cyclical communication's process data.

Due to this flexibility, safely connecting different parts of the machine becomes more simple. The safety container is routed through the various controllers and processed in the various parts of the machine. This makes emergency stop functions for an entire machine or bringing targeted parts of a machine to a standstill easily possible – even if the parts of the machine are coupled with other communication systems (e.g. Ethernet).

Implementing the FSoE protocol in a device requires little resources and can lead to a high level of performance and correspondingly, short reaction times. In the robotics industries, there are applications that use FSoE for safe motion control applications in an 8-kHz closed loop.



Black Channel principle: the standard communication interface can be used.


Further information regarding Safety over EtherCAT can be found on the ETG website: www.ethercat.org/safety

Communication profiles

In order to configure and diagnose slave devices, it is possible to access the variables provided for the network with the help of acyclic communication. This is based on a reliable mailbox protocol with an auto-recover function for erroneous messages.

In order to support a wide variety of devices and application layers, the following EtherCAT communication profiles have been established:

- CAN application protocol over EtherCAT (CoE)
- Servo drive profile, according to IEC 61800-7-204 (SoE)
- Ethernet over EtherCAT (EoE)
- File access over EtherCAT (FoE)
- ADS over EtherCAT (AoE)



Different communication profiles can coexist in the same system.

A slave device isn't required to support all communication profiles; instead, it may decide which profile is most suitable for its needs. The master device is notified which communication profiles have been implemented via the slave device description file.

CAN application protocol over EtherCAT (CoE)

With the CoE protocol, EtherCAT provides the same communication mechanisms as in CANopen®-Standard EN 50325-4: Object Dictionary, Mapping of Process Data Objects (PDO) and Service Data Objects (SDO). Even the network management is similar. This makes it possible to implement EtherCAT with minimal effort in devices that were previously outfitted with CANopen, and large portions of the CANopen Firmware are even reusable. Optionally, the legacy 8-byte PDO limitation can be waived, and it's also possible to use the enhanced bandwidth of EtherCAT to support the upload of the entire Object Dictionary. The device profiles, such as the drive profile CiA 402, can also be reused for EtherCAT.


Servo drive profile according to IEC 61800 7 204 (SoE)

SERCOS™ is known as a real-time communication interface, especially for motion control applications. The SERCOS™ profile for servo drives is included in the international standard IEC 61800-7. The standard also contains the mapping of this profile to EtherCAT. The service channel, including access to all drive-internal parameters and functions, are mapped to the EtherCAT mailbox.

Ethernet over EtherCAT (EoE)

EtherCAT makes use of physical layers of Ethernet and the Ethernet frame. The term Ethernet is also frequently associated with data transfer in IT applications, which are based on a TCP/IP connection.

Using the Ethernet over EtherCAT (EoE) protocol any Ethernet data traffic can be transported within an EtherCAT segment. Ethernet devices are connected to an EtherCAT seg-



Transparent transmission of standard IT protocols

ment via Switchports. The Ethernet frames are tunneled through the EtherCAT protocol, similarly to the internet protocols (e.g. TCP/IP, VPN, PPPoE (DSL), etc.) as such, which makes the EtherCAT network completely transparent for Ethernet devices. The device with the Switchport property takes care of inserting TCP/IP fragments into the EtherCAT traffic and therefore prevents the network's real-time properties from being affected.

Additionally, EtherCAT devices may also support Ethernet protocols (such as HTTP) and can therefore behave like a standard Ethernet node outside of the EtherCAT segment. The master device acts as a layer-2 switch that sends the frames via EoE to the corresponding nodes according to their MAC addresses. In this way, all internet technologies can also be implemented in an EtherCAT environment, such as an integrated web server, E-mail, FTP transfer, etc.

File access over EtherCAT (FoE)

This simple protocol similar to TFTP (Trivial File Transfer Protocol) enables file access in a device and a uniform firmware upload to devices across a network. The protocol has been deliberately specified in a lean manner, so that it can be supported by boot loader programs. A TCP/IP stack isn't required.

ADS over EtherCAT (AoE)

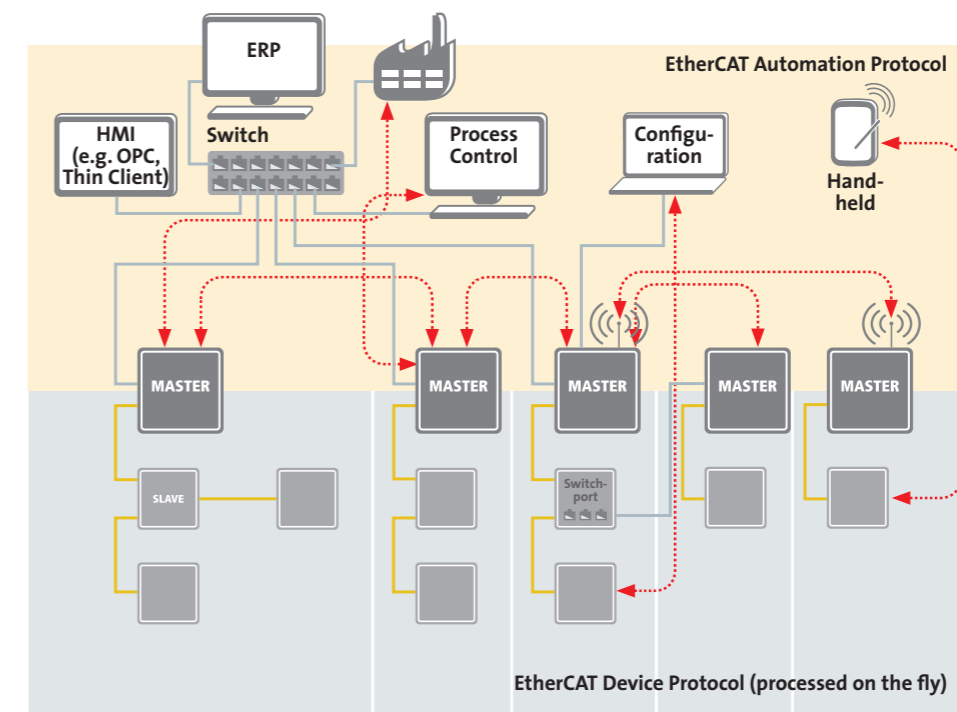
As a mailbox-based client-server protocol, ADS over EtherCAT (AoE) is defined by the EtherCAT specification. While protocols such as CAN application protocol over EtherCAT (CoE) provide a detailed semantic concept, AoE complements these perfectly via routable and parallel services wherever use cases require such features. For example, this might include access to sub-networks through EtherCAT using gateway devices from a PLC program such as CANopen®, IO-Link™ and others.

AoE comes with far less overhead when compared with similar services provided by the Internet Protocol (IP). Both sender and receiver addressing parameters are always contained in the AoE telegram. As a result, a very lean implementation on both ends (client and server) is possible. AoE is also the protocol of choice for acyclic communication via the EtherCAT Automation Protocol (EAP) and therefore provides seamless communication between an MES system, the EtherCAT master, and subordinated fieldbus devices connected via gateways. AoE serves as the standard means to obtain EtherCAT network diagnostic information from a remote diagnostics tool.

Plant-wide communication with the EtherCAT Automation Protocol (EAP)

The process control level has special communication requirements that differ slightly from the requirements addressed by the EtherCAT Device Protocol, described in the previous sections. Machines or sections of a machine often need to exchange status information and information about the following manufacturing steps with each other. Additionally, there is usually a central controller that monitors the entire manufacturing process, provides the user with status information on productivity, and assigns orders to the various machine stations.


The EtherCAT Automation Protocol (EAP) fulfills all of the above requirements.



Factory-wide communication with EtherCAT

The protocol defines interfaces and services for:

- Exchanging data between EtherCAT master devices (master-to-master communication)
- Communication to Human Machine Interfaces (HMI)
- A supervising controller to access devices belonging to underlying EtherCAT segments (routing)
- Integration of tools for the machine or plant configuration, as well as for device configuration



Factory-wide communication architecture with the EtherCAT Automation Protocol and Safety over EtherCAT

The communication protocols used in EAP are part of the international standard IEC 61158. EAP can be transmitted via any Ethernet connection, including a wireless link, for example making it possible to include automated guided vehicles (AGV), which are common in the semiconductor and automotive industries.


Cyclic process data exchange with EAP follows either the push or poll principle. In push mode, each node sends its data either with its own cycle time or in a multiple of the own cycle time. Each receiver can be configured to receive data from specific senders. Configuring the sender and receiver data is done through the familiar object dictionary. In poll mode, a node (often the central controller) sends a telegram to the other nodes, and each node responds with its own telegram.

The cyclic EAP communication can be directly embedded within the Ethernet frame, without additional transport or routing protocol. Again, the EtherType 0x88A4 identifies the EtherCAT-specific use of the frame. This enables the exchange of high-performance data with EAP in a millisecond cycle. If data routing between distributed machines is required, the process data can also be transmitted via UDP/IP or TCP/IP.

Additionally, with the help of the Safety over EtherCAT Protocol, it's also possible to transmit safety-critical data via EAP. This is common in cases where parts of a large machine need to exchange safety-critical data to realize a global emergency stop function, or to inform neighboring machines of an emergency stop.

Integration of other bus systems

EtherCAT's ample bandwidth makes it possible to embed conventional fieldbus networks as an underlying system via an EtherCAT gateway, which is particularly helpful when migrating from a conventional fieldbus to EtherCAT. The changeover to EtherCAT is gradual, making it possible to continue using automation components that don't yet support an EtherCAT interface.



Decentralized fieldbus interfaces

The ability to integrate decentralized gateways also reduces the physical size of the Industrial PC, sometimes even to an embedded Industrial PC, since extension slots are no longer necessary. In the past, extension slots were also required to connect complex devices, such as fieldbus master and slave gateways, fast serial interfaces, and other communication subsystems. In EtherCAT, all that is needed to connect these devices is a single Ethernet port. The process data from the underlying subsystem is made directly available in the process image of the EtherCAT system.

EtherCAT

Powering the digital transformation with EtherCAT, Industrie 4.0 and IoT



EtherCAT

Process optimizations, predictive maintenance, manufacturing as a service, adaptive systems, resource-savings, smart factories, cost reductions: there are countless good reasons to utilize control network data in higher level systems. The Internet of Things (IoT), Industrie 4.0, Made in China 2025, Industrial Value Chain Initiative – there is a common need across the board for seamless, continuous and standardized communication across all levels. Sensor data uploaded into the cloud along with recipes and parameters downloaded from ERP systems into distributed devices; take, for example, a feeding system shared by two machines: there are data flow requirements both in vertical and horizontal directions.

EtherCAT inherently meets the requirements of the digital transformation through its high performance, flexibility and open interfaces:

Superior system performance is the prerequisite for adding big data features to control networks.

EtherCAT provides the flexibility to add cloud connectivity to existing systems without even touching the controller or updating the slave devices: Edge gateway can access any data within any EtherCAT slave device via the mailbox gateway feature of the EtherCAT master. The edge gateway can either be a remote device, talking to the master via TCP or UDP/IP, or a software entity directly located on the same hardware as the EtherCAT master itself.

Additionally, open interfaces allow one to integrate any IT-based protocol – including OPC UA, MQTT, AMQP or any others – either within the master or directly into the slave devices, thus providing a direct link for IoT without protocol discontinuities from the sensor to the cloud.

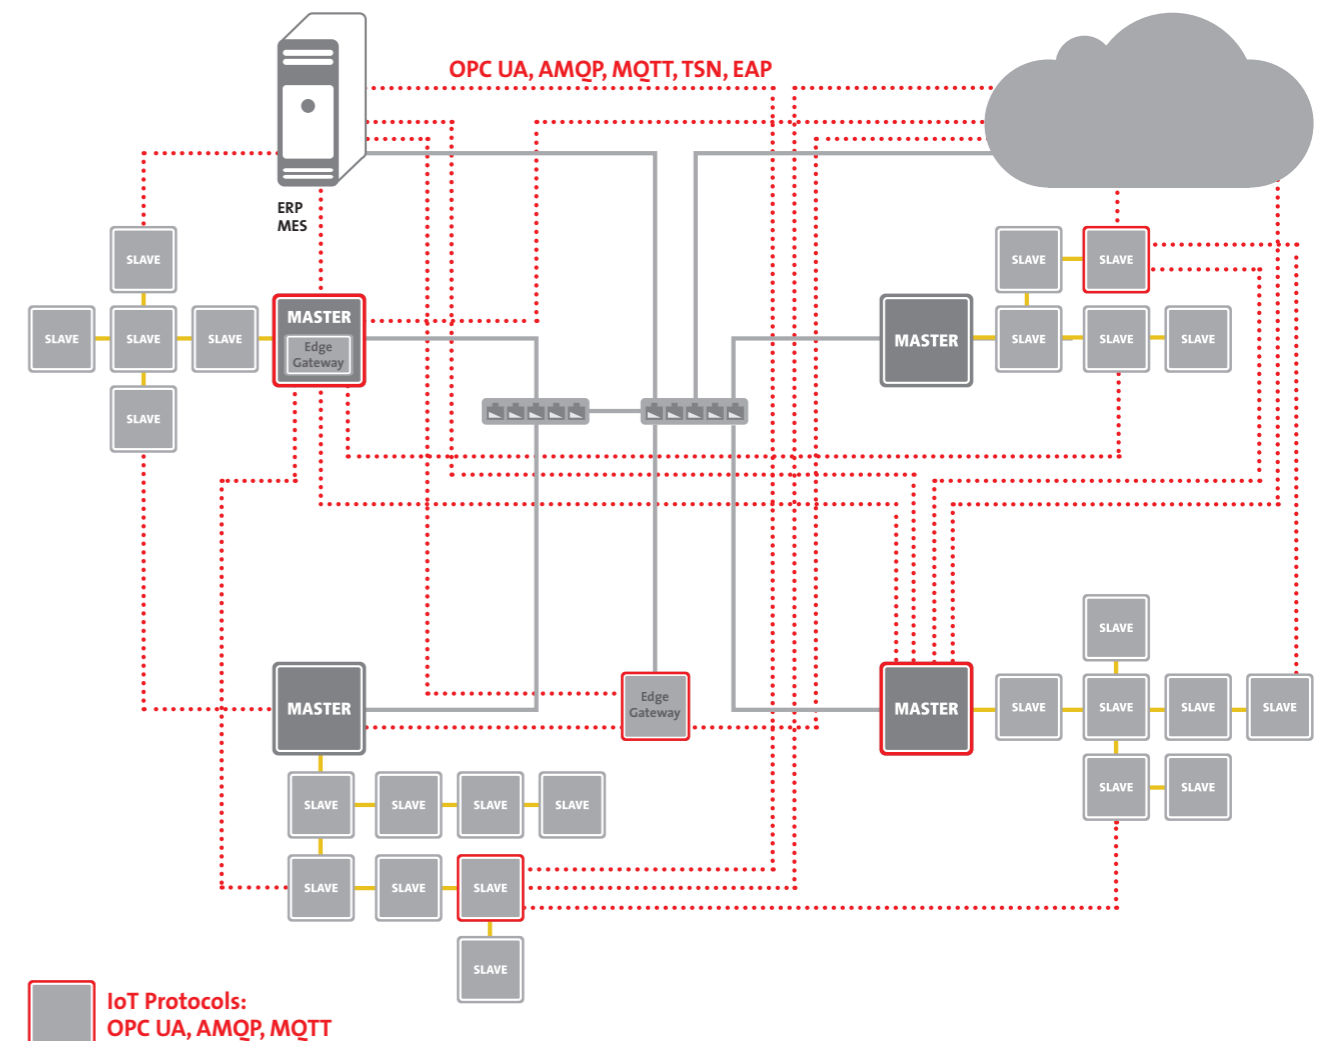
All these features have always been part of the EtherCAT protocol, which shows how forward-thinking the architecture is. Nevertheless, more networking features are added as they evolve and become relevant. Of course, it is also important to consider the past when looking forward: this introduction of valuable new features is managed with total network continuity as the EtherCAT protocol itself has been stable at “Version 1” since its introduction in 2003.

Other new developments in the area of Time-Sensitive Networking (TSN) features further improve the real-time capabilities of the controller-to-controller communication. Enabled by TSN, control systems – even those that are cloud-based – can access a network of EtherCAT slaves also across plant networks. Since EtherCAT typically only needs one frame for an entire network, such access is much leaner and thus faster than any other fieldbus or industrial Ethernet technology. In fact, EtherCAT Technology Group experts have contributed to the TSN task group of IEEE 802.1 from day one – at a time when TSN was still known as AVB.

Further information about EtherCAT and TSN can be found here: www.ethercat.org/tsn

The EtherCAT Technology Group (ETG) was also among the first fieldbus organizations to partner with the OPC Foundation. The OPC UA protocol complements the EtherCAT portfolio because it is a scalable TCP/IP based client/server communication technology with integrated security, enabling encrypted data transfer up to MES/ERP systems.

With OPC UA Pub/Sub, the usability of OPC UA has been improved in machine-to-machine (M2M) applications and for vertical communication to cloud based services. The ETG is actively contributing to all these developments to ensure that they fit seamlessly into the EtherCAT environment. Therefore, EtherCAT is not only IoT-ready, EtherCAT is IoT!






Implementing EtherCAT interfaces

EtherCAT technology has been specially optimized to enable low-cost design, so adding an EtherCAT interface to a sensor, I/O device, or embedded controller should not significantly increase device costs. Furthermore, the EtherCAT interface also doesn't require a more powerful CPU; the CPU requirements instead are based only on the needs of the target application.

In addition to hardware and software requirements, development support and the availability of communication stacks are important when developing an interface. The EtherCAT Technology Group offers worldwide developer support by quickly answering questions or addressing technical issues. Finally, evaluation kits available from multiple manufacturers, developer workshops, as well as free sample code make getting started a little easier.

For the end user, the most important factor is the interoperability of EtherCAT devices from various manufacturers. To ensure interoperability, device manufacturers are required to perform a conformance test prior to bringing their device on the market. The test checks if the implementation follows the EtherCAT specification, and can be performed with the EtherCAT Conformance Test Tool. The test can also be used during device development to discover and correct implementation issues early on.



The interface for an EtherCAT master device has a single, unbelievably simple, hardware requirement: an Ethernet port. The implementation uses either the on-board Ethernet controller or an inexpensive standard network card, so no special interface card is required. That means that with just a standard Ethernet port, a master device can implement a hard real-time network solution.

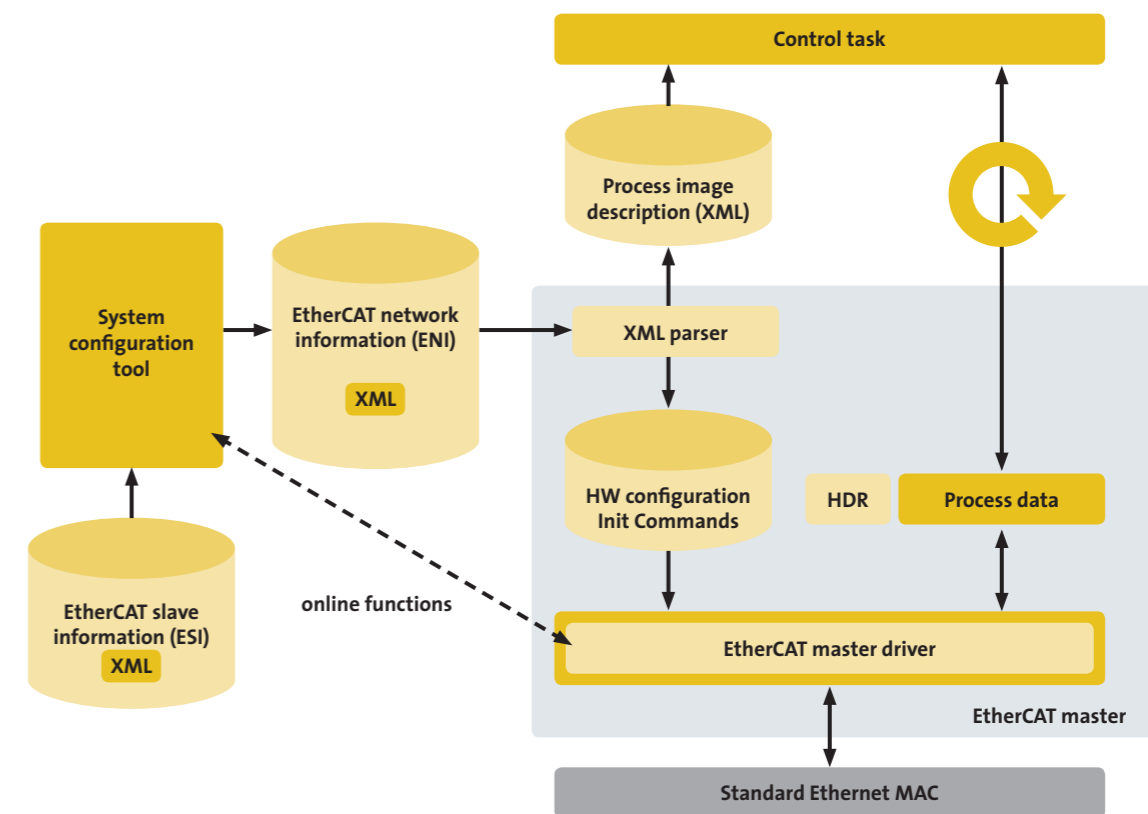
In most cases the Ethernet controller is integrated via Direct Memory Access (DMA), so no CPU capacity is required for the data transfer between the master device and the network. In an EtherCAT network, mapping occurs at the slave devices. Each slave device writes its data to the right location in the process image and reads the data addressed to it all while the frame is moving through. Therefore, the process image that arrives at the master device is already sorted correctly.

Since the master device CPU is no longer responsible for the sorting, its performance requirements depend only on the desired application and not the EtherCAT interface. Especially for small, mid-sized, and clearly defined applications, implementing an EtherCAT master is a snap. EtherCAT master devices have been implemented for a wide variety of operating systems: Windows and Linux in various iterations, QNX, RTX, VxWorks, Intime, eCos are just a few examples.

ETG members offer a variety of options to support the implementation of an EtherCAT master, ranging from free downloads of EtherCAT master libraries, sample master code, all the way to complete packages (including services) for various real-time operating systems and CPUs.

In order to operate a network, the EtherCAT master requires the cyclic process data structure as well as boot-up commands for each slave device. These commands can be exported to an EtherCAT Network Information (ENI) file with the help of an EtherCAT configuration tool, which uses the EtherCAT Slave Information (ESI) files from the connected devices.

The breadth of available master implementations and their supported functions varies. Depending on the target application, optional functions are either supported or purposely omitted to optimize the utilization of hardware and software resources. For this reason, EtherCAT master devices are categorized into two classes: a Class-A-Master is a standard EtherCAT master device, while a Class-B-Master is a master device with fewer functions. In principle, all master implementations should aim for Class A classification. Class B is only recommended for cases in which the available resources are insufficient to support all functionalities, such as in embedded systems.



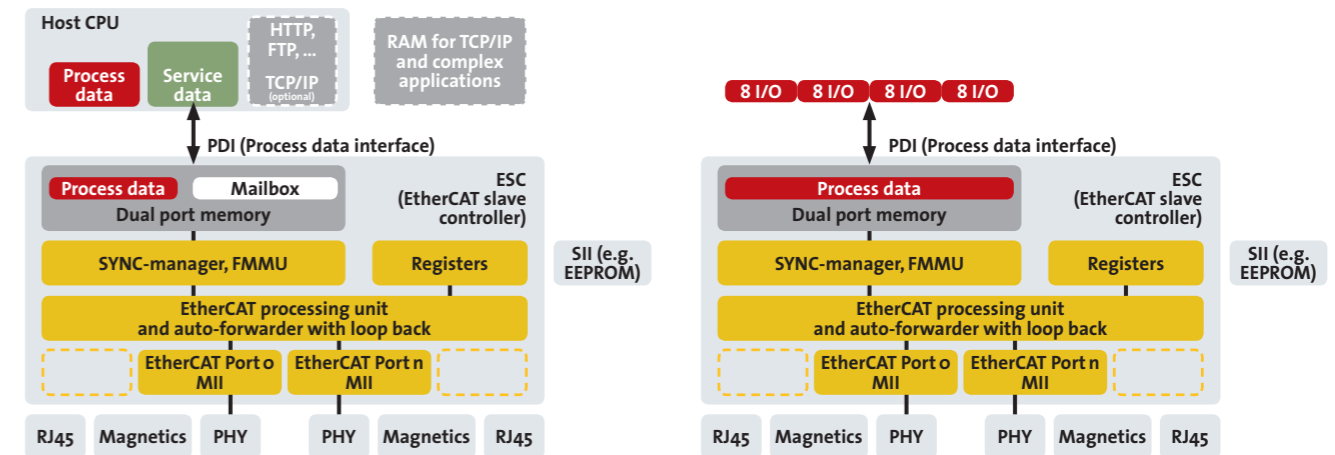

Typical EtherCAT master architecture

Implementing a slave device

EtherCAT slave devices use inexpensive EtherCAT Slave Controllers (ESC) in the form of an ASIC, FPGA, or integrated in a standard microcontroller. Simple slave devices don't even need an additional microcontroller, because inputs and outputs can be directly connected to the ESC. For more complex slave devices, the communication performance depends only minimally on the microcontroller performance, and in most cases, a 8-bit microcontroller is sufficient.

EtherCAT Slave Controllers are available from multiple manufacturers, with the size of the internal DPRAM and the number of Fieldbus Memory Management Units (FMMUs) depending on the variation. Different Process Data Interfaces (PDI) for external access from the application controller to the application memory are also available:

- The 32-bit parallel I/O Interface is suitable for connecting up to 32 digital inputs and outputs, but also for simple sensors or actuators for which 32 data bits are sufficient and no application controller is required.
- The Serial Peripheral Interface (SPI) is targeted at applications with small amounts of process data, such as analog I/O devices, encoders, or simple drives.
- The parallel 8/16-bit microcontroller interface corresponds to common interfaces of fieldbus controllers with integrated DPRAM. It is particularly suitable for complex nodes with larger amounts of data.
- Suitable synchronous busses for various microcontrollers have been implemented for FPGA and System-on-Chip (SoC) variations.



Slave hardware:
EtherCAT slave controller with host CPU

Slave hardware:
EtherCAT slave controller with direct I/O

The hardware configuration is stored in a non-volatile memory (e.g. an EEPROM), the Slave Information Interface (SII), which contains information about the basic device features, so that the master can read this at boot-up and operate the device even if the device description file is not available. The EtherCAT Slave Information (ESI) file that comes with the device is XML based and contains the complete description of its network accessible properties, such as process data and their mapping options, the supported mailbox protocols including optional features, as well as the supported modes of synchronization. An EtherCAT network configuration tool uses this information for online and offline configuration of the network.

Various manufacturers offer evaluation kits for implementing slave devices. These kits include slave application software in source code, and they sometimes also include a test master. Using an evaluation kit, a fully functional master slave EtherCAT network can be commissioned in just a few steps.

The ETG website contains the ETG.2200 Slave Implementation Guide with useful tips and hints on further documentation for implementing slave devices:

www.ethercat.org/etg2200



Two of the most important factors for a communication standard to be successful are conformance and interoperability. That's why the EtherCAT Technology Group takes both of these factors very seriously. In addition to requiring a conformance test for each device implementation (aided by the automated EtherCAT Conformance Test Tool), the ETG offers a wide variety of activities to ensure the interoperability between EtherCAT master devices, slave devices, and also EtherCAT configuration tools.

Plug Fests

When trying to test if multiple devices are interoperable, one of the most pragmatic things to do is to try connecting the devices together. With this approach in mind, the ETG holds multiple Plug Fests each year, with each Plug Fest usually spanning two days. During the Plug Fests, master and slave device manufacturers come together to test how their devices behave together, which improves the usability of devices in the field. Members can exchange EtherCAT tips and tricks and have their questions answered by EtherCAT experts. The ETG holds Plug Fests in Europe, North America and Asia.

EtherCAT Conformance Test Tool

The EtherCAT Conformance Test Tool (CTT) makes it possible to automatically test an EtherCAT slave device's behavior.

It's a Windows program requiring only a standard Ethernet port. The tool sends EtherCAT frames to the Device under Test (DuT) and receives its responses. A test case is marked as passed if the response from the DuT corresponds to the defined response. The test cases are defined as XML files. This makes it possible to modify or expand the test cases without having to modify the actual test tool. The TWG Conformance is responsible for specifying and releasing the most current valid test cases.

In addition to the protocol tests, the CTT also examines if the values in the EtherCAT Slave Information (ESI) file are valid. Finally, the CTT performs device-specific protocol tests, such as for the CiA 402 drive profile.

All of the testing steps and results are saved in a test logger, and can be analyzed or saved as a documented verification for the device release.

The ETG consistently maintains and adds new test cases to the CTT. It is important that a device manufacturer always has the most recent version of the tool for testing products prior to release. To make it easier, the CTT is offered as a subscription. The CTT is also useful during the design phase to uncover early errors in the interface implementation.

Technical Working Group Conformance

The EtherCAT Conformance Test Policy requires device manufacturers to test each device with a valid version of the EtherCAT Conformance Test Tool before the device is brought on the market. The manufacturer may conduct this test in-house.

The ETG Technical Committee (TC) established a Technical Working Group (TWG) Conformance, which determines the test procedures, the contents of the test, and the implementation of the Conformance Test Tool. The TWG Conformance is continuously expanding the tests and their depth.

The TWG Conformance also established an interoperability test process, with which devices can be tested in the context of an entire network.

EtherCAT Test Centers

The official EtherCAT Test Centers (ETC) in Europe, Asia and North America are accredited by the ETG and perform the official EtherCAT Conformance Test. The EtherCAT Conformance Test includes the automated tests run with the CTT, interoperability tests within a network, as well as an examination of the device's indicators, markings, and tests of the EtherCAT interfaces.

Device manufacturers are encouraged but not obligated to have their devices tested at an ETC. After the conformance test is passed, the manufacturer receives an EtherCAT Conformance Tested certificate for the device. This certificate is only available for devices that have passed the official conformance test at an ETC; not for those which have been tested in-house.

The additional test in an accredited EtherCAT Test Center further increases compatibility and the uniform operation and diagnostics of EtherCAT implementations. End users should be sure to ask for the EtherCAT Conformance Tested certificate when choosing devices for their application.

More information on conformance and certification is available on the ETG website: www.ethercat.org/conformance

www.ethercat.org

The EtherCAT website provides comprehensive information about the technology as well as upcoming events, the latest EtherCAT products and the current membership directory. Also, there are focus topics such as functional safety and conformance of EtherCAT devices. Furthermore, the website provides presentations, press articles and publications in the download section.

EtherCAT Product Guide

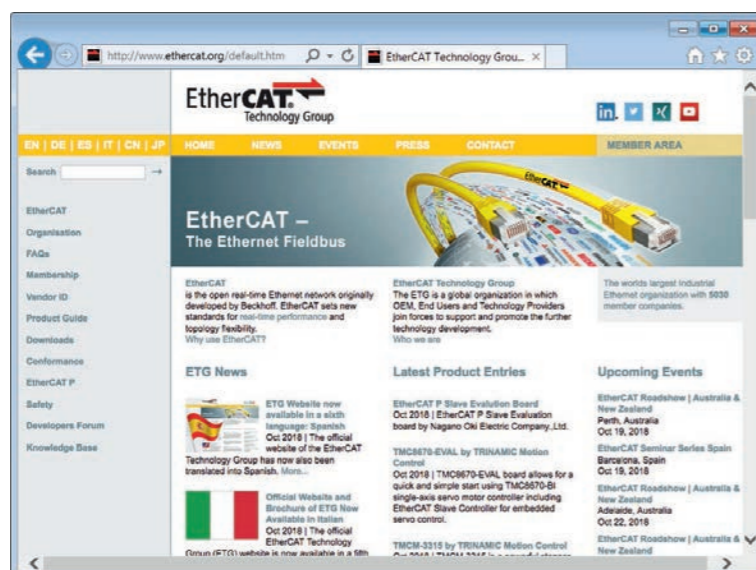
The EtherCAT Product Guide is a directory listing EtherCAT products and services based on information provided by ETG members, and is available online at www.ethercat.org/products. If you have any questions about the products, please contact the manufacturer directly, as the ETG itself does not sell any products.

Event Section

The event section shows the worldwide events hosted by the ETG and those that are organized in conjunction with the association. In the calendar located at www.ethercat.org/events important dates can be found, including those for the technical working group meetings, trade show appearances, EtherCAT workshops and Industrial Ethernet seminars.

Member Area

Members have insider access to the protected area of the website at www.ethercat.org/memberarea, which contains valuable additional items such as all EtherCAT specifications, the online developer forum and a knowledge base with all necessary information for implementation, configuration and diagnosis of EtherCAT devices and networks.



ETG worldwide



Contact

ETG Headquarters

Ostendstraße 196
90482 Nuremberg
Germany
Phone: +49 (911) 5405620
Fax: +49 (911) 5405629
info@ethercat.org

ETG Office North America

Carlsbad, CA, USA
Phone: +1 (877) ETHERCAT
info.na@ethercat.org

ETG Office China


Beijing, P. R. China
Phone: +86 (0)10 82200090
Fax: +86 (0)10 82200039
info@ethercat.org.cn

ETG Office Japan

Yokohama, Japan
Phone: +81 (45) 6501610
Fax: +81 (45) 6501613
info.jp@ethercat.org

ETG Office Korea

Seoul, Korea
Phone: +82 (0)2 2107 3242
Fax: +82 (0)2 2107 3969
info.kr@ethercat.org



EtherCAT®, EtherCAT P®, Safety over EtherCAT®, are trademarks or registered trademarks, licensed by Beckhoff Automation GmbH, Germany. Other designations used in this publication may be trademarks whose use by third parties for their own purposes could violate the rights of the owners.