White Paper

Built-in Cybersecurity: EtherCAT Delivers Proven Protection

Dr. Guido Beckmann[1], Technical Committee Chair, EtherCAT Technology Group
Florian Essler[1], Conformance TWG Chair, EtherCAT Technology Group
Torsten Förder[1], Cybersecurity Expert, Beckhoff Automation
Alexander W. Koehler[2], S&S Principal Security Advisor for Cybersecurity, UL Solutions Germany
Fatma Kotzaoglan, Senior Project Engineer, UL Solutions United Kingdom
Martin Rostan, Executive Director, EtherCAT Technology Group

[1]Member of CENELEC TC65X WG3 Cyber Security
[2]DKE TBKON Chair of the Cybersecurity Advisory Board for Conformity Assessment

EtherCAT provides cybersecurity capabilities and system security enablement which meet the cybersecurity requirements for Security Level 2 in accordance with IEC 62443. For nearly all common applications, no modification or extension of the EtherCAT protocol is necessary for CRA compliance. For attacks up to Security Level 3, backward-compatible software extensions are sufficient. This is confirmed by UL Solutions who assessed three types of systems based on the EtherCAT field bus and certified those systems under IEC 62443-3-3.

The European Cyber Resilience Act (CRA) has placed cybersecurity on the agenda for fieldbus systems as well. The CRA defines cybersecurity requirements for products with digital elements sold in Europe; this includes not only automation components but also the machines and systems equipped with them. Similar laws are being drafted in Asia and in the US, making the issue relevant worldwide.

In the future, therefore, fieldbus systems must also meet the CRA's requirements. The industry is addressing this challenge in very different ways.

Traditional fieldbus systems, as well as some Industrial Ethernet systems, are falling back on a kind of cell protection concept—essentially "IT zoning" in the OT domain: Communication from the cell or zone to the outside is protected by IT security measures such as firewalls, encryption, and access controls, so that no additional protection is required for communication and components within the cell. However, effective protection requires that no negligent or targeted cyberattack can occur within the cell: Consequently, operators and other personnel must be effectively prevented from gaining direct access to the fieldbus.

This may work if the fieldbus system is confined to a locked control cabinet. However, if the cell spans an entire machine or plant where fieldbus cables are physically accessible, this becomes difficult. Zoning also contradicts the "from sensor to cloud" claim—a selling point of some bus architectures that becomes obsolete due to security requirements.

Other approaches require integrating IT security into every field device. With authenticated encryption extending all the way to the I/O module, the CRA's requirements can also be met for Industrial Ethernet systems that are accessible to operating personnel. However, this not only noticeably degrades the performance of the bus system but also increases complexity: Managing certificates that must be regularly renewed not only complicates commissioning but, above all, significantly complicates the maintenance of machines and systems.

Since EtherCAT uses standard Ethernet frames (IEEE 802.3) and standard Ethernet physical layers, it is sometimes assumed that, from a cybersecurity perspective, it must be treated like an Ethernet network and protected using the same proven IT measures as those for switch-based Ethernet networks.

However, this is only partially true, as the unique characteristics of EtherCAT make many attack scenarios familiar from the IT world impossible from the outset, rendering the standard defensive measures unnecessary. If this is not taken into account, IT security experts will check whether the standard IT measures are in place—and the absence of such measures will generally result in a downgrade. However, this does no more justice to the unique characteristics of EtherCAT than requiring a diesel particulate filter would do for an electric car.

EtherCAT takes a different approach, based on IEC 62443, the international standard for the cybersecurity of industrial automation and control systems. IEC 62443 is considered the internationally recognized standard in this field and is currently being expanded in Europe to serve as the basis for assessing systems and components in accordance with the CRA. The standard is also a good fit for the CRA because both follow a risk-based approach: as many cybersecurity measures as necessary, not as many as possible. In doing so, risks must always be assessed in the context of the specific application and the expected severity of attacks: An industrial washing machine is significantly less at risk than an energy distribution system that is part of critical infrastructure.

IEC 62443-3-3 describes the security requirements for *systems* (not for individual components) and is therefore the appropriate standard for a system such as EtherCAT.

The standard breaks down the seven foundational requirements into over 100 system requirements and uses security levels to categorize the strength of attacks to be considered.

Security Level (SL) 1 corresponds to protection against accidental or unintentional disruptions, such as those caused by malware or ransomware in the factory network. SL 2 covers targeted attacks by actors without specialized expertise. SL 3 already assumes the attacker has specialized expertise, corresponding resources, and motivation. Finally, SL 4 addresses attackers with capabilities and resources typically available only to state actors—effective protection against such attacks is practically unachievable and, in most cases, economically unjustifiable.

How EtherCAT Works

A basic understanding of how EtherCAT works is necessary to grasp the technology's unique characteristics and, consequently, its specific cybersecurity features. Here is a brief overview:

EtherCAT is an Industrial Ethernet technology in which the so-called MainDevice—typically the automation system controller—sends Ethernet frames that are processed by the SubDevices, i.e., the field devices. The MainDevice sends the frames via a conventional Ethernet interface with a Medium Access Controller (MAC) in accordance with IEEE 802.3. Instead of the MAC, the SubDevices use a special EtherCAT SubDevice Controller (ESC). The ESC is not capable of sending frames itself.

EtherCAT utilizes the principle of processing on-the-fly. Unlike conventional Industrial Ethernet networks, in which a MAC receives the Ethernet frame intended for the field device and then forwards it to the host controller's software stack for processing, EtherCAT continuously processes the frames in the ESC. The EtherCAT SubDevice therefore forwards the frame while it is still being received. The frame experiences only the minimum physical delay in the process. The principle of on-the-fly processing requires the hardware implementation of the ESC functionality: Since only a few bits of the frame are accessible at any given time, only a few nanoseconds are available for processing. The ESC functionality therefore cannot be implemented in software. From a cybersecurity perspective, the hardware implementation of the EtherCAT SubDevice Controller has the advantage that it can only be manipulated with great effort.

In addition, the frames are always processed with a fixed delay. Some of the chip's most important functions are also inaccessible to the application controller in the field device: For example, the setting that determines which data is

to be extracted from or inserted into the passing Ethernet frame can only be changed externally via the EtherCAT bus system—that is, by the MainDevice—but not by potentially faulty or manipulated firmware in the field device.

EtherCAT System Properties

Let us now consider the specific system-related properties of EtherCAT and their implications for potential cybersecurity attacks:

1. The EtherCAT system-architecture minimizes the attack surface. No software stack has access to the frames as they pass through the SubDevice; only the hardware logic processes them. Only EtherCAT frames are present in the medium. Only the MainDevice can generate EtherCAT frames.

2. The EtherCAT protocol is directly embedded in the Ethernet frame and is thus located at Layer 2 of the ISO-OSI model. The Internet Protocol (IP) at Layer 3 is not used for transporting EtherCAT payload data. Since the primary function of IP is to route data across subnets, malware is practically always dependent on the presence of the IP layer: Malicious content embedded directly within the Ethernet frame would be unable to cross subnet boundaries. This means that EtherCAT field devices are inherently protected against common malware attacks.

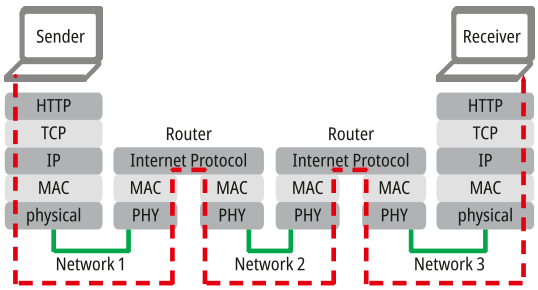

IP as a Prerequisite for Malware Routing

EtherCAT uses its own EtherType (0x88A4), followed by the EtherCAT header. This allows EtherCAT SubDevice Controllers to immediately recognize, based on the frame and the EtherCAT header, whether the incoming frame is an EtherCAT frame.

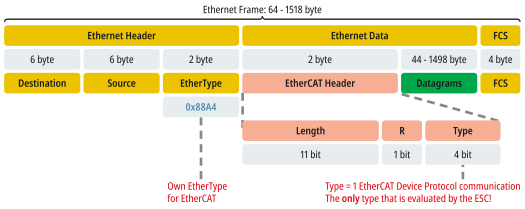

EtherCAT: Directly in the Ethernet Frame

All other frame types are discarded by the EtherCAT SubDevice Controller hardware: A non-EtherCAT frame arriving in the EtherCAT segment is thus reliably removed by the next field device. This means that attacks via non-EtherCAT frames are impossible from the outset.

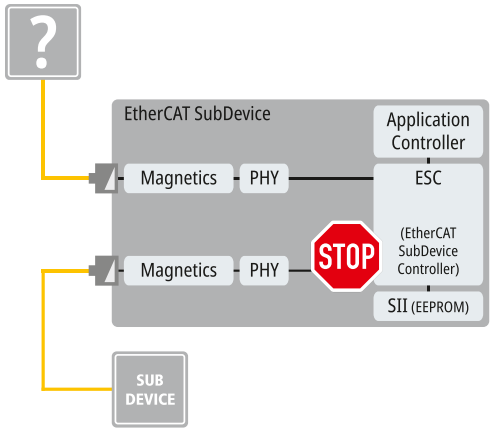

ESC destroys all non-EtherCAT frames

3. EtherCAT SubDevices cannot modify data that is not addressed to them. This applies equally to process data and parameter data. Since the FMMU—the hardware unit in the ESC responsible for mapping process data between local and global memory—cannot be configured via the SubDevice's local ESC interface, firmware on the field device that has been accidentally or intentionally manipulated cannot activate it on the bus. Thus, an attack involving the insertion of a software-manipulated field device has no effect on other devices.

4. When the network starts up, the EtherCAT MainDevice can query the properties of all connected SubDevices and detect the

network topology. The level of detail of this query can be adapted to the application's requirements: For example, the MainDevice can limit itself to querying only the device type and its communication properties at startup, or it can also compare the firmware version and serial number, thereby detecting not only the addition of devices and changes to the topology but also the replacement of devices.

5. EtherCAT SubDevice Controllers have up to 4 ports and therefore support a wide variety of network topologies. Unused ports can be disabled by the MainDevice and thus blocked in the hardware.

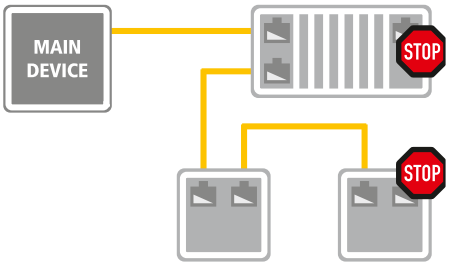

Disabling Unused Ports

In this case, connecting a device via an unused port does not result in a logical connection: Although a link signal is established, communication with the ESC is blocked. Thus, an attack by plugging a device into an open port is ineffective and can even be detected via the additional link.

6. Data traffic in EtherCAT networks can be monitored and analyzed using so-called "probes." Probes are hardware devices that must be installed in the network between the MainDevice and field devices or between two field devices. However, these devices cannot manipulate data—they can only receive it. A probe can therefore be used to record data traffic, but not to modify data. It is extremely complex to deduce the function of the machine or system from the content of the recorded process data, and without precise knowledge of the machine mechanics and system properties that are not transmitted via the bus (e.g., naming of process variables), this is nearly impossible. So, if the goal of the attack is to understand how the machine or system works, an attack on the control system is almost always more promising and effective.

7. Even though EtherCAT itself does not use the Internet Protocol (IP), the technology is nevertheless capable of tunneling any Ethernet traffic using the "Ethernet over EtherCAT" extension, thereby, for example, retrieving data from web servers on field devices, transporting OPC UA or other protocols, including IP-based traffic. Since this data stream is tunneled, its content is completely independent of the functionality of EtherCAT itself and all devices involved in the transport. In principle, however, it can compromise the security of inadequately protected field devices at the end of the tunnel—i.e., the dedicated recipient of the data stream from the tunnel. If this is an attack vector that must be considered and mitigated, the "Ethernet over EtherCAT" protocol can be disabled either globally or selectively within the MainDevice.

8. The MainDevices can write the configuration data to the SubDevices at every startup. This is possible thanks to the high efficiency of EtherCAT. As a result, any configurations that may have been written to the device offline are overwritten, rendering this attack vector ineffective as well.

9. Regarding Denial-of-Service (DoS) attacks: Since incoming frames are processed by the ESC's hardware, bursts are not a problem: EtherCAT frames are processed immediately. A three-buffer mechanism in the ESC ensures that memory is always available for incoming process data, even if the device firmware cannot keep up with emptying the buffers. Acyclic data such as parameters are acknowledged, so it is not possible for the protocol stack to become overloaded. Non-EtherCAT frames are immediately discarded

by the ESC hardware. Thus, DoS attacks are ineffective.

10. As shown, due to the characteristics of the EtherCAT SubDevice Controller, a SubDevice can neither modify data nor send frames independently. Inserting a SubDevice is therefore not a viable attack scenario; furthermore, in EtherCAT, the addition of a device can be detected via the resulting change in signal propagation time. ESC chips measure frame propagation delay with an accuracy of well under 50 nanoseconds. The EtherCAT MainDevice reads these values to ensure high-precision synchronization of the distributed clocks. A significant deviation in these values—such as after a reboot—indicates a change in the network topology, as would occur when a device is inserted.

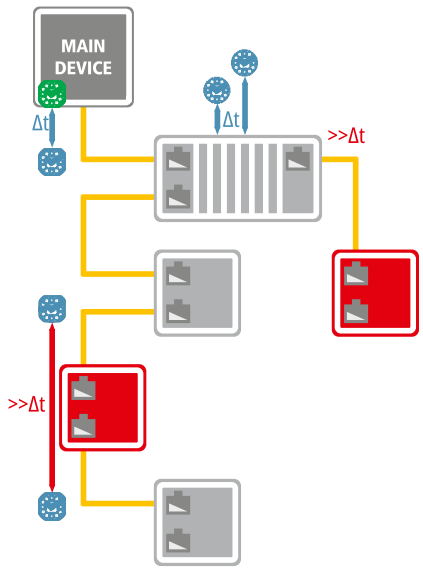

Inserted devices are detected via EtherCAT runtime measurement

11. Even the addition of a SubDevice equipped with “malicious” EtherCAT SubDevice Controller hardware developed specifically for this purpose is detected in this way. Developing such a chip requires considerable specialized expertise and significant effort, making this attack vector highly inefficient. Only in very few systems is it conceivable that the achievable effect—in this case, presumably the sabotage of the system—would economically justify the effort, or that it would not be simpler and more cost-effective to achieve this via other, non-EtherCAT-based attack scenarios.

And what about the control system?

In addition to EtherCAT, most EtherCAT control systems have additional interfaces, such as an Ethernet port for connecting to higher-level networks, memory card slots, or USB ports. EtherCAT field devices may also have additional interfaces of a similar nature. These must, of course, be very carefully protected against cybersecurity attacks through appropriate measures! This essential part of the cybersecurity strategy is unaffected by the EtherCAT network and therefore cannot be included in the security considerations for EtherCAT.

EtherCAT MainDevice implementations range, by design and intentionally, from very simple devices for low-demand applications to high-performance controllers that support significantly more EtherCAT features. Some security features of EtherCAT also depend on the implementation of the MainDevice. For example, not all controllers check the network for changes during the startup process or disable unused ports on SubDevice devices by default. To meet the relevant cybersecurity requirements, it is therefore necessary to select suitable MainDevices or controllers.

The EtherCAT Technology Group develops profile specifications that guide users, system integrators, machine builders, and device manufacturers on which measures must be used or implemented to achieve an appropriate security level.

EtherCAT achieves Security Level 2 without modifications

Due to the system characteristics mentioned above, EtherCAT technology already meets the cybersecurity requirements for systems expected to face attacks of Security Level 2 strength, even without security-specific

technology extensions. This was demonstrated in a comprehensive analysis: Each of the more than 100 system requirements of IEC 62443-3-3 was examined and evaluated in detail for a typical EtherCAT plant network. UL Solutions Germany has assessed the security capabilities of the EtherCAT protocol in operation in three typical application scenarios which lead to three certificates IEC 62443-3-3, based on security capabilities' requirements, selected in consequence of threat scenarios as outlined above. Those three certificates have been published by UL Solutions (Demko), Denmark, which is the Accepted National Certification Body by IECEE with recognition in the EU as well as globally because of IECEE's global MRA network.

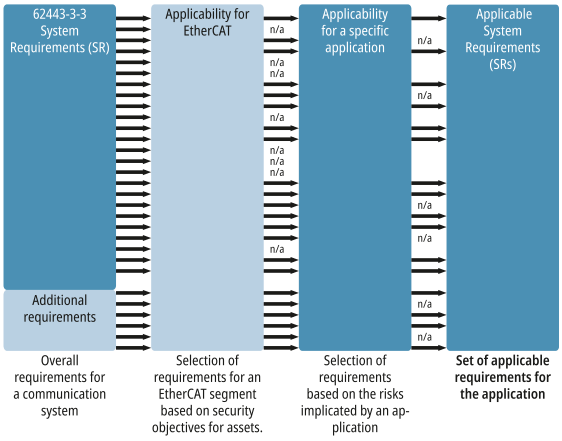

System Requirements: Assessment of applicability and degree of compliance for EtherCAT in the scenarios under consideration.

Higher security levels without hardware changes—software only

For systems requiring protection against an attacker's strength of Security Level 3, software extensions in the MainDevice and in the configuration tool are sufficient. This has also been confirmed by UL Solutions. The SubDevices can therefore remain unchanged. Among other things, runtime measurement is activated for protection against this attack strength, which detects additional devices in the EtherCAT segment as described above.

The cybersecurity risk associated with manipulated device description files is manageable in EtherCAT: after all, these are read only once by the tool that creates the network configuration.

A modified file could lead to altered device behavior, which would likely be noticed during system commissioning or functional testing. If not: The EtherCAT Technology Group (ETG) is establishing its own Certificate Authority (CA) so that ETG members can easily and uniformly sign and authenticate EtherCAT device description files and software.

Optional encryption for EtherCAT without certificate handling

Only in very few applications do users additionally require encryption of the process and parameter data transmitted via the bus system to protect it from eavesdropping: while "encryption" appears on many cybersecurity checklists, in practice the damage—and thus the risk—that could result from data disclosure is usually minimal. For systems with exceptionally high cybersecurity requirements, the ETG is extending the EtherCAT standard—with full backward compatibility—to include a quantum-secure, authenticated encryption method that secures all or only a portion of the data, as desired. Key distribution occurs without certificates, eliminating the need to renew certificates due to expiration dates and significantly simplifying management. "Fully backward-compatible" even means that the ESC chips do not need to be modified: This functionality can be added purely via software—in both the MainDevice and the SubDevice.

References:

UL Solutions: 3 Certificates IEC 62443-3-3 and Technical Reports TRF: Solution Application of Capabilities Assessment of EtherCAT Technology. issued by UL Solutions (Demko) Denmark, IECEE Certification Body: DK-177530-UL/DK-178394-UL/DK-178399-UL. IECEE CBTL (testing lab): UL Solutions Northbrook, IL, USA.