



Simplifying Rail Metro Ethernet Communications Networks

R. Weatherburn – Vertical Marketing, Transportation
C. Burns - Field Application Engineering Manager, Transportation
Dr. O. Kleineberg – Advanced Development

Executive Summary

To give passengers access to the internet and other wireless services onboard and to both manage and control increasingly sophisticated trains, the rail metro communications network philosophy is changing. Operators are increasingly making Ethernet the bearer of choice. This move to Ethernet-based networks comes at a time when capacity demands on rail metro systems are greatly increasing due to global trends that are driving the public transportation market. These trends are:

- Urbanization
- Concern over greenhouse gases and global warming
- Reducing carbon resources
- De-regulation
- Increased Digitalization – The Rail IoT

This white paper briefly describes two of these trends and the resulting multiple trackside Ethernet networks. IEEE Time Sensitive Networks (TSN) technology is introduced and then its use (along with redundancy protocols) is proposed to enable radical reduction in the complexity of rail metro Ethernet networks, hence reducing both CAPEX and OPEX costs.

Table of Contents

Executive Summary	1
Trends Driving Multiple Trackside Networks	2
Increasing Urbanization Driving Metro Capacity Demands	2
Increasing Digitization and Passenger Expectations Driving Ethernet Network Deployment	3
Multiple Trackside Ethernet Networks – Conclusion	4
TSN – Time Sensitive Networking	5
TSN - Mechanisms and Interdependencies	5
Prioritization Based on Timing with the Time-Aware Scheduler	5
Necessity of Guard Bands and Interruption of Ethernet Frames	6
Synchronous Transmission Cycles are a Prerequisite	7
Preventing Interfering Traffic with Ingress Filtering and Policing	8
Implementing Rail Metro Networks with TSN	9
Summary and Outlook	10

Be certain.
Belden.

Trends Driving Multiple Trackside Networks

Increasing Urbanization Driving Metro Capacity Demands

The trend of increasing urbanization coupled with climate change, environmental concerns and increasing urban mobility mean that the demands on urban railway system capacity have never been greater.

Traditionally, increasing the capacity on an urban rail network has involved large levels of capital expenditure on the rail infrastructure. Typical solutions to capacity challenges may have involved laying additional running lines, perhaps boring new tunnels and lengthening platforms to enable longer trains. These large civil projects are not only expensive, they are also potentially disruptive to running an operational service.

Through the use of exact, real-time, train positioning information, Communications Based Train Control (CBTC) provides a solution to the capacity challenges and avoids the need for many of these expensive and disruptive civil engineering capital projects. An increased utilization of the existing rail infrastructure is realized through the removal of traditional fixed signaling blocks and instead managing train headways as well as speed and acceleration profiles continuously in real time. A CBTC system managing shorter, lighter, faster trains running more closely together will result in greater capacity vs. longer trains running on a fixed block signaling system.

A large proportion of the total cost of running rail passenger vehicles is fixed for long durations (such as vehicle

leasing or finance costs). Fuel or energy usage very often represents the largest proportion of costs that the operator can control. The use of a CBTC system allows energy savings to be made through the application of different driving profiles (such as coasting and reduced acceleration curves) when passenger demand is less; at off-peak travelling times, for example. When coupled with passenger counting technology, this management of the rail vehicles can be made fully automatic.

CBTC systems run over a digital networked radio system by means of antennas or radiating cable for the bi-directional communication between the track equipment and the trains. The 2.4GHz band is commonly used in these systems (public WiFi band), though other alternative frequencies such as 900 MHz (US), 5.8 GHz or other licensed bands may be used as well (the use of licensed bands is becoming more common to ensure that interference is less likely).

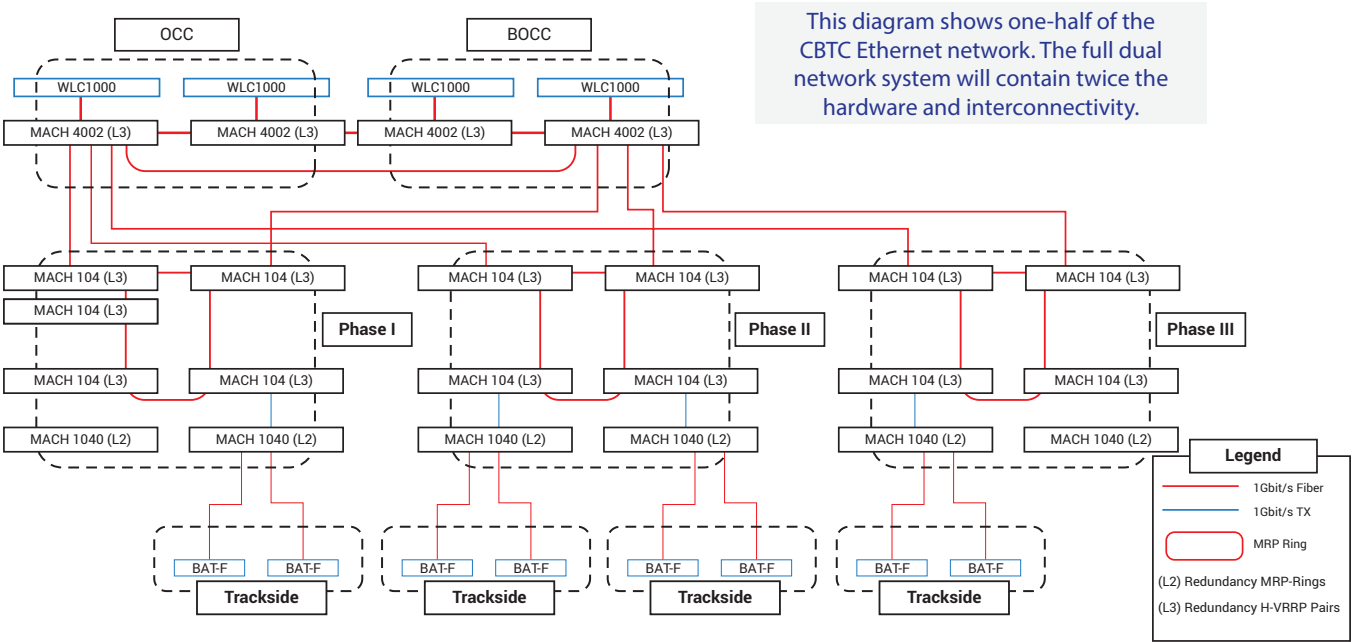
Some of the key requirements of the CBTC communications network are shown in the table below.

Requirements	Benefits
Appropriate Railway Onboard Vehicle Certification <ul style="list-style-type: none"> • IEC/EN for Europe • AAR/AREMA for North America 	A pre-requisite for installation onboard train builders' vehicles or trackside and for fitment by system integrators. Without this certification it is very difficult to sell equipment into the rail industry.
Fast, Secure Roaming	For safety reasons, any disruption in the CBTC service will bring trains to a stop. Because of this, a roaming handover time of <50ms is essential to provide uninterrupted train-to-ground communication.
Error Tolerance	While the communications system for CBTC is not defined as safety critical, loss of the data being transmitted will result in system errors. A packet loss of less than 0.1% is required throughout the system.
Sufficient Bandwidth	The communications system must be capable of carrying all of the required data from trains to the Operational Control Centre in order for the CBTC system to operate reliably. To provide this service, a bandwidth of at least 4Mbps is required throughout the communications network.
Network Latency	Since the CBTC operation relies on real-time information for the control of moving trains, the maximum latency from end to end must be less than 5ms.

To comply with these requirements, the communications network being deployed must be designed for the exclusive use of the CBTC system. Typically the communications network is made up of two interlaced trackside Ethernet networks, with the rail vehicle being in constant wireless connection with at least one of the two networks. Wireless access points for each network are spaced and placed down the length of the rail infrastructure. This interlaced architecture helps overcome any wireless handover delays between access points and the moving vehicle. In addition, the CBTC communications network is 'ring fenced' to guarantee the bandwidth and latency performance criteria. This means that CBTC Ethernet networks exclusively carry data relating to the safety-critical CBTC application.

Figure 1 shows an example of an Ethernet network architecture deployed for CBTC. Note that the network diagram only shows one of the interleaved networks!

Figure 1 - One half of a typical CBTC Ethernet Communications Network



Increasing Digitization and Passenger Expectations Driving Ethernet Network Deployment

As our requirements for trains become more demanding, the diagnostic and control systems supporting them also need to grow. Some of the demands include:

- increased fuel efficiency
- increased reliability
- shorter dwell times at platforms
- reduced headway between trains
- fewer onboard staff (perhaps no driver)
- improved security

To understand onboard equipment reliability and to ensure that the train performs as expected, we monitor and collect data on a whole range of onboard systems. Increasingly, the data flow does not stop at just being onboard; diagnostic systems now report back to depots before the train arrives for maintenance so that spares and tools can be organized. Passenger Information Systems provide us with real-time information while ticketing, seat reservation systems and WiFi access to the internet or entertainment systems make our journey less stressful.

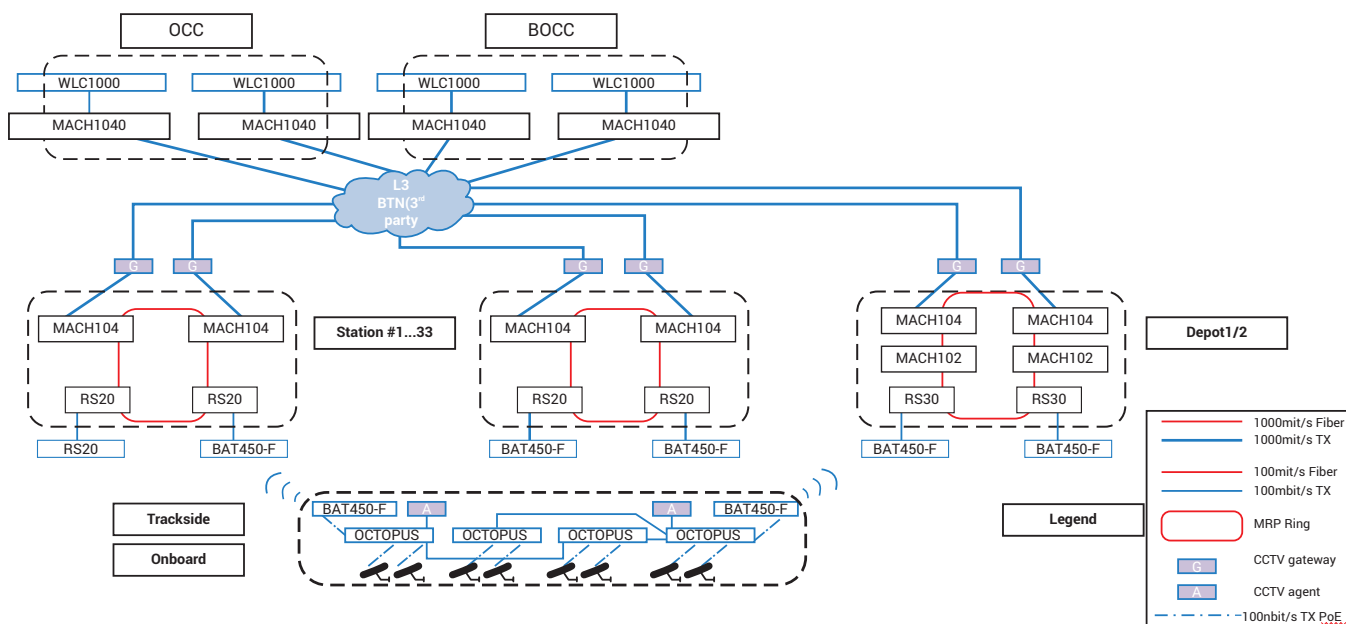
The number of Ethernet based applications that need a train-to-ground connection to function is on the increase; Figure 2 shows the most common applications currently deployed around the world (grouped by application primary function).

Figure 2 - Commonly deployed applications requiring train-to-ground connectivity

Safety	Reliable & Punctual	Productive & Entertaining
Surveillance	Remote Diagnostics	Passenger WiFi / Internet
Platform Edge Systems	Platform Edge Systems	Entertainment
Driver Only Video	Passenger Counting	Seat Reservation Systems
Help / Panic Points	Passenger Information System	Passenger Information Systems
Wireless Signaling	Wireless Signalling	
Front Facing & Pantograph Video	Front Facing & Pantograph Video	
Fire Protection Systems	Driver Assistance System	
	Ticketless Travel	

These applications require a trackside wireless infrastructure and the ability to transfer the data to the remote control room or monitoring workstation. As we noted previously, this data cannot be carried on the CBTC Ethernet Communications infrastructure as the signaling application will no longer have the guaranteed network performance needed. This means that more Ethernet network infrastructure is required, with Figure 3 showing an example of a deployed network for a real-time video surveillance system.

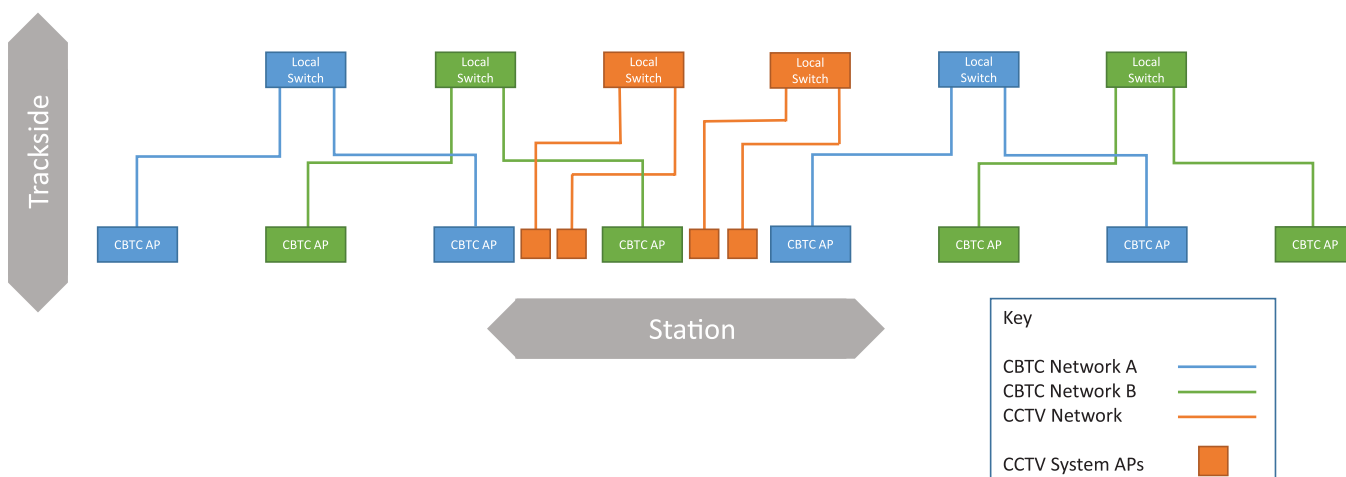
Figure 3 - Real-Time Video Surveillance Ethernet Network



Multiple Trackside Ethernet Networks - Conclusion

The deployment of CBTC and other trackside networks means that there is a large amount of Ethernet equipment installed trackside to backhaul the associated data to the Control or Operational centers. Figure 4 shows a simplified network diagram for the trackside architecture. To put the amount of hardware being installed in our two examples into perspective, the typical spacing of CBTC wireless APs is 200m to 300m (dependent on tunnel topography). Within a city metro system, stations are, on average, 1km apart (often less).

Figure 4 - Overall Trackside Network Architecture



This means that for a 30 km stretch of metro railway, there could be 240 wireless APs alone, made up as follows:

- 30 stations with 4 APs each
- 120 devices
- CBTC Network A (500m apart)
- 60 devices
- CBTC Network B (500m apart)
- 60 devices

When the associated cabling (copper and fiber), trackside switches and backbone data backhaul equipment are also considered, it's obvious that the current methodology for network deployment in metro systems could be made more efficient. Considerable cost savings could also be achieved (both during deployment and through the life of the systems).

Now, let's explore Time Sensitive Networking (TSN) and its potential to achieve these network deployment efficiencies.

TSN – Time Sensitive Networking

Time-Sensitive Networking (TSN) is currently being developed at the IEEE as a novel technology that offers an entirely new level of determinism in standard IEEE 802.1 and IEEE 802.3 Ethernet networks. This means that future communications networks in the transportation industry will be able to provide:

- Calculable, guaranteed end-to-end latencies for traffic of different priorities on the same network
- Highly limited latency fluctuations (jitter)
- Extremely low packet loss.

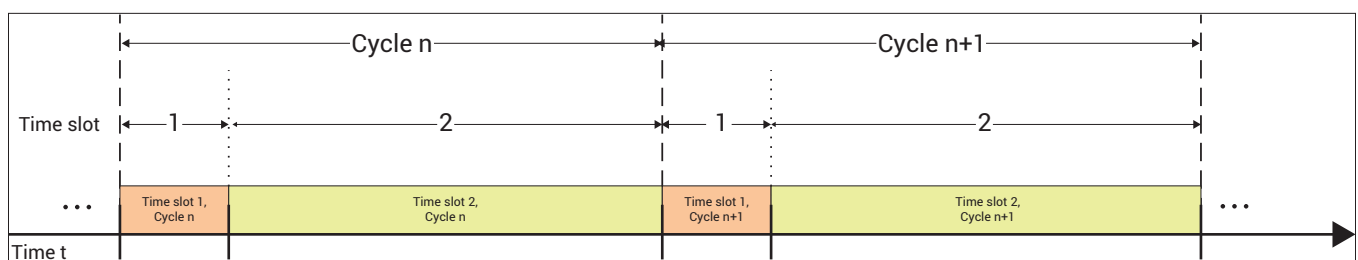
TSN - Mechanisms and Interdependencies

TSN adds a level of determinism to Ethernet-based data communication that is able to meet even the highest demands placed on an Ethernet network. TSN has been conceived as a modular system by which the precise characteristics of the implementation (and the associated hardware and software requirements) can be tailored to fit the individual network requirements. Appropriately so, TSN is not made up of a single standard document, but it is a family of standards which have been in development by the IEEE 802.1 TSN Task Group¹ since 2012. These activities have yielded their first results: central mechanisms of the TSN family are already available as standard documents. In order to give an overview of these new technologies, the following sections will address the most important TSN mechanisms and their interdependence.

Prioritization Based on Timing with the Time-Aware Scheduler

Until now, it was not possible with the strict priorities of Class of Service (CoS) mechanisms such as the IEEE 802.1Q to guarantee bounded end-to-end latency of time-sensitive data traffic. Due to queueing effects, an Ethernet frame with low priority that is already in transmission could delay Ethernet frames of even the highest priority (7) at every Ethernet switch along the transmission path. As one of the central components of TSN, the Time-Aware Scheduler (TAS) introduces the possibility to assign specific time slices for transmission based on the CoS Ethernet frame priority. Thus, traffic of a specific priority is granted exclusive access to the transmission medium. Forwarding and delivery of the Ethernet frames is guaranteed at a defined point in time.

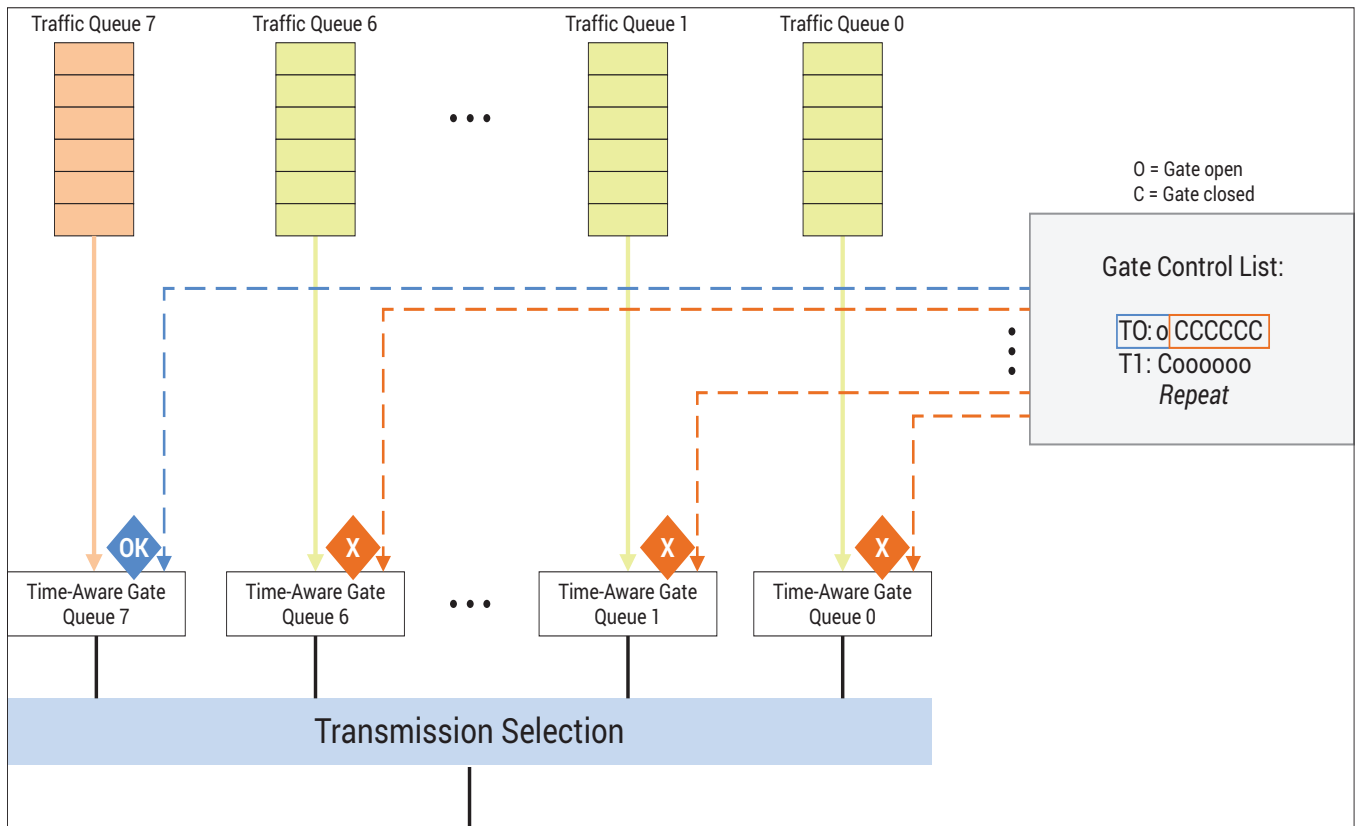
Figure 5 - Time division multiplexing permits the reservation of time-slots within a cycle in order to enable the timely transmission of periodic real-time data



The fundamental idea of this TSN mechanism -- published as Standard IEEE 802.1Qbv-2016² in March 2016 -- is to utilize TDMA (Time Division Multiple Access) to divide time into discreet segments of equal length, so-called cycles, as illustrated in Figure 5. This allows dedicated time slots to be provided for the transmission of data packets with real-time requirements within the cycles. With the aid of the Time-

Aware Scheduler, the transmission of conventional Best Effort Ethernet traffic can be temporarily interrupted in order to forward time-sensitive data traffic within the reserved time slots for high-priority traffic. The Time-Aware Scheduler thus permits the prioritization of periodic real-time data (see Time slot 1 in Figure 5) in relation to conventional Best Effort data traffic.

Figure 6 - The Time-Aware Scheduler implements the time-based prioritization



Similar to the strict prioritization scheme, the Time-Aware Scheduler utilizes the CoS priorities (PCP - Priority Code Point) that are present in the VLAN tag of the Ethernet header. In this case, all Ethernet frames are processed until they reach the Time-Aware gate queues at the output port. At this point, the Time-Aware Scheduler intervenes in the packet processing, as illustrated in Figure 5. More precisely, with the use of the Time-Aware Scheduler, the selection of the next Ethernet frame to be transmitted is no longer just determined strictly by a hierarchy at the queue, but rather the state of the respective gates is also taken into consideration. This state may be either open or closed, based on actual time. Ethernet frames that are waiting for transmission in the associated queues will be considered in the packet selection, depending on these states. In Figure 6, for example, only the queue with a priority of 7 is processed at this particular point in time.

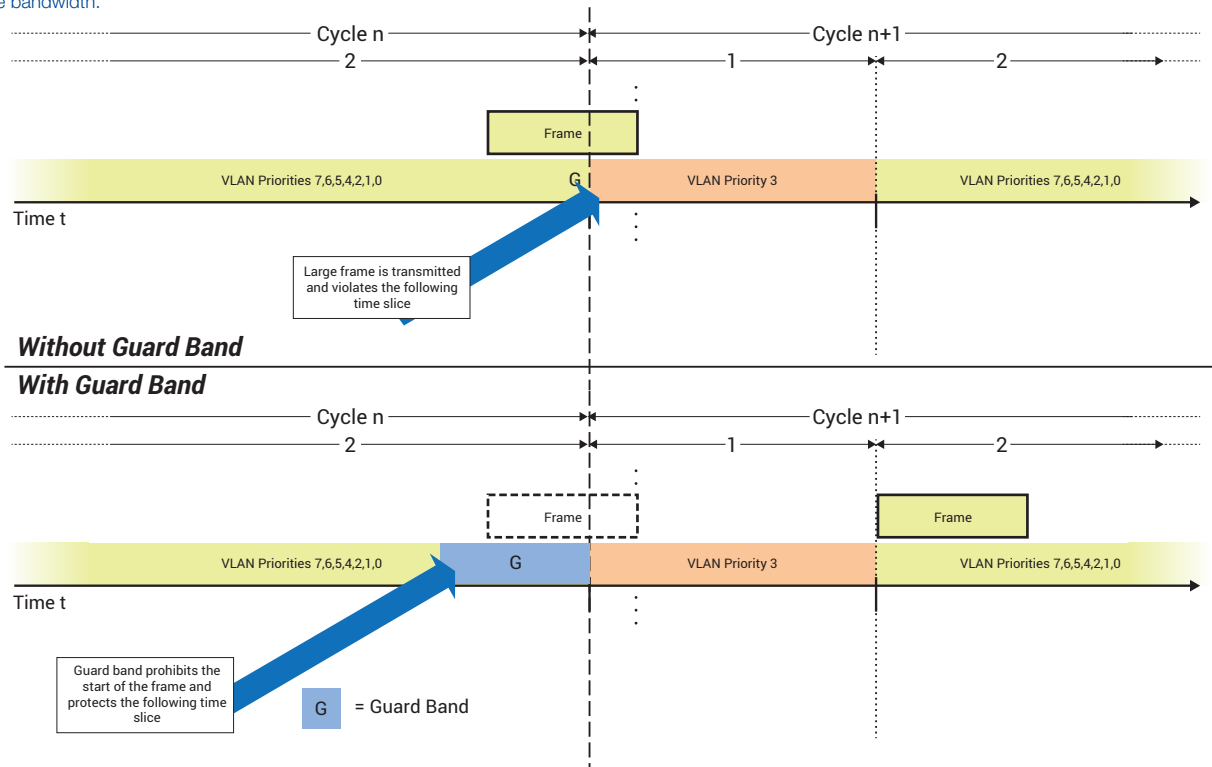
The Gate Control List determines which traffic queue is permitted to transmit at a specific point in time within the

cycle. Besides the states of the Time-Aware Gates, the Gate Control List indicates the length of time for which a specific entry will be active. In case of the Gate Control List shown on the right side in Figure 6, the list mirrors the cycle that consists of a Best Effort phase, as well as a phase with prioritized data traffic from Figure 5.

Necessity of Guard Bands and Interruption of Ethernet Frames

Due to the very poor predictability of Best Effort traffic patterns, it is generally not foreseeable when a specific Best Effort data packet will need to be processed. As illustrated in Figure 7, for example, the transmission of an Ethernet frame in time slot 2 could be initiated too late. This Ethernet frame would then, despite the use of the Time-Aware Scheduler, extend into the time slot number 1 of the subsequent cycle. This would therefore result in a delayed processing of real-time data and a violation of guaranteed end-to-end latencies.

Figure 7 - The guard band in TSN prevents Best Effort frames from extending into a time slot that is reserved for real-time data, but it decreases the available bandwidth.



In order to maximize the usable bandwidth for Best Effort Ethernet frames, the IEEE 802 working group developed a method for Ethernet frame pre-emption (IEEE 802.1Qbu3, IEEE 802.3br4), which was completed in June 2016. With this method, conventional Ethernet frames can be divided into partial packets ("framelets") of as small as 64 bytes and each framelet may be transmitted separately. As shown in Figure 5, this permits starting the transmission of a large Ethernet frame, despite insufficient remaining time within the Best Effort phase. The frame can be interrupted at the last 64 byte boundary before the current time slot ends and can then be completed in the next Best Effort phase. Frame pre-emption makes it possible to reduce the guard band to the maximum size of one Ethernet framelet plus the maximum remaining transmission time (63 byte times) that cannot be pre-empted

any more. In the case of a fast Ethernet network, for example, the dead time from each guard band can be reduced to ~0.12 ms and thus, a significant improvement of the use of the bandwidth available can be achieved.

Due to the fact that frame pre-emption is a significant intrusion into the normal process of Ethernet frame forwarding and processing, it is necessary for both devices of an Ethernet connection (e.g. two Ethernet switches) to announce their support for this mechanism through the use of the Link Layer Discovery Protocol (LLDP) identified in IEEE 802.1AB5). Only with frame pre-emption support on both ends of the link can the feature be activated on the corresponding end devices or switch ports. With this, backwards compatibility with existing Ethernet devices is maintained.

Figure 8 - With the method of Ethernet frame pre-emption, the guard band size can be reduced from the maximum size of an Ethernet frame to the size of a partial packet.

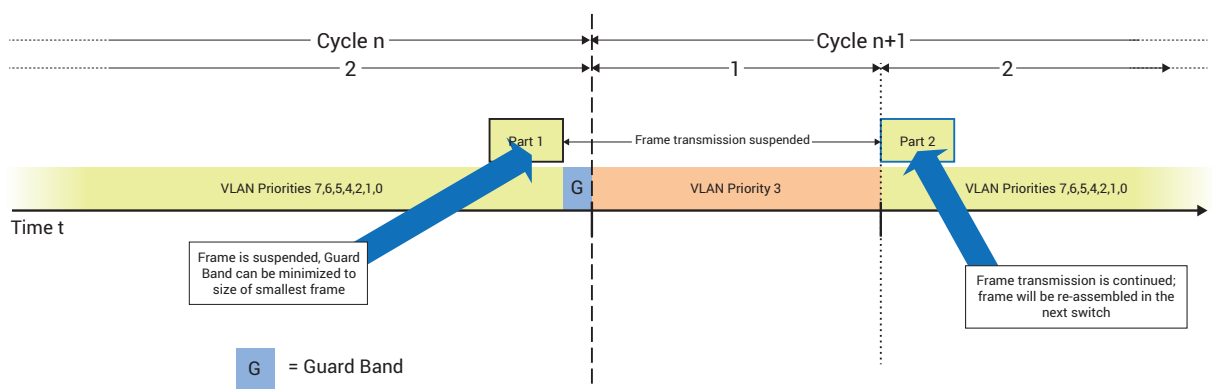
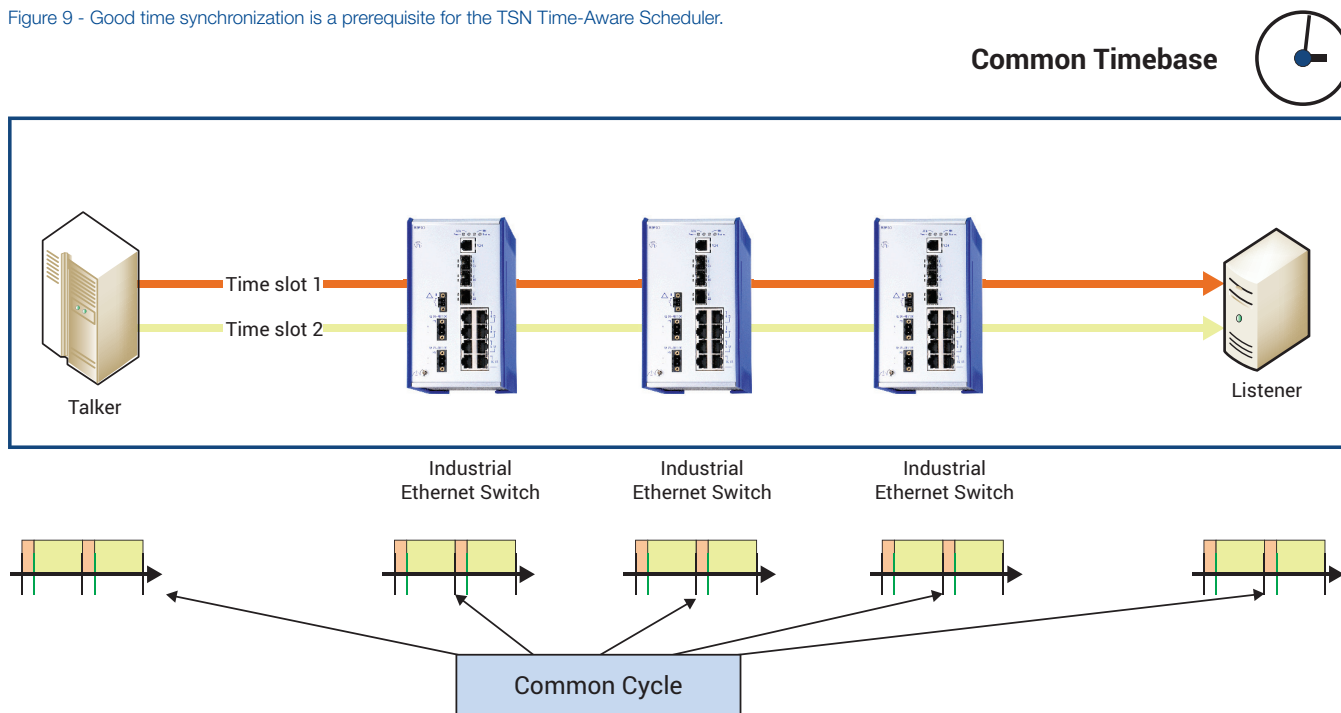


Figure 9 - Good time synchronization is a prerequisite for the TSN Time-Aware Scheduler.



Synchronous Transmission Cycles are a Prerequisite

The Time-Aware Scheduler utilizes only local configuration data- the data that is available in a particular network device (end device or switch). This configuration data consists of information about the lengths of cycles and time slots, for example. Therefore, besides the Time-Aware Scheduler, close coordination between the devices in the network is required in order to ensure that the frames match the correct time slots in each switch. This enables the transmission of communication streams that can be transmitted through end-to-end connections, with guaranteed latencies and without queuing times (see Figure 9). This means, in particular, that all network participants must possess a common understanding of time. Specifically, all participants must know when a cycle begins and which time slot is active in the cycle. In order to enable this, the use of a protocol for time synchronization, such as the Precision Time Protocol (PTP) in accordance with IEEE 1588 or any available IEEE 1588 profile, such as IEEE 802.1AS7 is mandatory.

Both IEEE 1588 as well as IEEE 802.1AS permit the synchronization of distributed clocks within a network with an accuracy of under 1 μ s. Implemented in hardware, timing precision in the range of a few nanoseconds can be achieved (Hirschmann PTP white paper8). In contrast to the protocols known from IT environments, such as Network Time Protocol (NTP), IEEE 1588 doesn't necessarily have to utilize a global synchronization with, for example, an atomic clock. More commonly, the network participant with the most precise, freely running clock is determined with the aid of the Best Master Clock (BMC) algorithm. This device then serves as the reference clock (Grandmaster Clock) against which all remaining network participants are synchronized. For TSN, it is of primary importance that the time is synchronized to all clocks within a network. The actual time of day, on the other hand, plays only a secondary role.

The IEEE 1588 profile, IEEE 802.1AS, follows the same fundamental synchronization model as PTP. It was originally developed to limit the large number of configuration options to those parameters that are relevant in local networks (LANs). For example, in case of the transport technology and encapsulation, IEEE 802.1AS is confined to Ethernet transport, while IEEE 1588 provides an additional IP (Internet Protocol) encapsulation scheme for use in wide area networks (WANs). As a result of the TSN standardization process, the existing IEEE 802.1AS profile has been expanded by the addition of parameters from IEEE 1588 that are required for use in automation networks. For example, IEEE 1588 offers support for multiple time domains that can be synchronized in parallel. Accordingly, with IEEE 1588, network participants can be synchronized with a global time reference (as with NTP), as well as a second network time reference. That offers the option to use the global synchronization for unambiguous event logging, while the network-wide synchronized clock can be used for the Time-Aware Scheduler, since in this case, synchronization according to global conventions such as the leap second is not required. Among other things, this capability will also be included in the next version of this profile with IEEE 802.1AS-Rev9.

Since the current version of IEEE 1588 was specified back in 2008, this technology for time synchronization has already been established in many markets and application areas. In some cases, profiles for special applications, such as the energy market, have been developed and are in use today. In these cases, there is no need to specifically utilize IEEE 802.1AS for time synchronization - the TSN mechanisms permit the use of any arbitrary mechanism for time synchronization. Thus, depending on the application area, IEEE 1588 can be used instead of IEEE 802.1AS, with or without a specific profile. In the future, IEEE 802 does not intend to limit this freedom of choice in regard to which protocol needs to be used for time synchronization. In any case, regardless of the synchronization protocol that is used, the quality of clock synchronization that is achieved is very high in order for all devices in the network to start-and-end cycles and time slices at the correct points in time.

Preventing Interfering Traffic with Ingress Filtering and Policing

In a system in which all participants behave as expected, the TSN standards previously described offer all of the mechanisms required for deterministic data transmission. However, the mechanisms discussed so far require a complete reception of frames, as well as (partial) frame processing in a forwarding switch or receiving end device. As a result, misconfigured devices or malicious network participants can significantly interfere with the operation of TSN mechanisms such as the Time-Aware Scheduler by sending data frames with erroneously assigned CoS priorities or by excessively stressing the resources assigned to them.

In order to counter this, an additional TSN mechanism is currently being developed within the IEEE 802.1 working group which allows discarding data frames that have been erroneously assigned at the time of reception (IEEE P802.1Qci10). In addition, this mechanism allows discarding real-time data streams that use more than their reserved bandwidth, thus allowing the policing of streams. Finally, TSN can make use of already existing Layer 2 security mechanisms, such as MACsec (IEEE 802.1AE11). This allows ensuring the authenticity of the sender so that only verified Ethernet frames are forwarded. This way, it is possible to handle a multitude of attacks and scenarios with erroneously configured network participants.

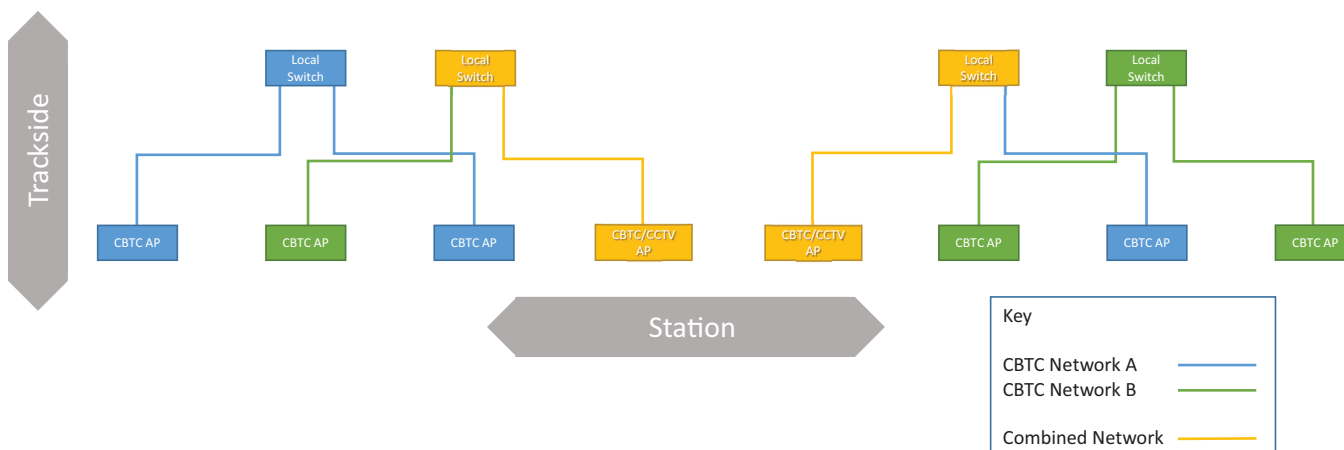


Implementing Rail Metro Networks with TSN

With TSN, deterministic data transmission with standardized Ethernet according to IEEE 802.1 and 802.3 is possible for the first time. Referring back to Figure 4, if we use TSN to combine the train-to-ground CCTV network present in the station with the CBTC network, we will simplify the network to that shown in Figure 10.

Using TSN to guarantee the transmission latency of the CBTC data allows us to utilize the CBTC network to carry the CCTV data from the station system, which reduces the number of trackside APs needed on our example 30km length of metro line by 60 devices. This is a significant saving when all of the associated cabling and network switches that are no longer required are also taken into account; by using TSN to guarantee the delivery of the CBTC data, all of the Ethernet equipment shown in Figure 3 is no longer required. In addition, the costs associated with using leased lines for data backhaul are no longer required, as the network installed and operated by the rail operator is now being utilized.

Figure 10 - Combined CBTC and Application (CCTV) Networks



Summary and Outlook

The standardization process in the area of time-sensitive networking is not yet completed and is expected to take a few more years. Accordingly, there are various TSN mechanisms that are currently still in the active standardization process. It is equally imaginable that additional mechanisms will be added to the already existing TSN family in the future.

Central mechanisms of the TSN protocol family have, however, been completed and have been demonstrated successfully. These mechanisms, such as the Time-Aware Scheduler, can already be integrated in products and their benefits can be used immediately. Equally, through the IEEE 802 standardization process, backwards compatibility is ensured: TSN networks that are already installed can still be used in the future. Therefore, TSN is no longer a future technology – its time has come.

References

- [1] <http://www.ieee802.org/1/pages/tsn.html>
- [2] <http://standards.ieee.org/findstds/standard/802.1Qbv-2015.html>
- [3] <http://standards.ieee.org/findstds/standard/802.1Qbu-2016.html>
- [4] <http://standards.ieee.org/findstds/standard/802.3br-2016.html>
- [5] <http://standards.ieee.org/findstds/standard/802.1AB-2016.html>
- [6] <https://standards.ieee.org/findstds/standard/1588-2008.html>
- [7] <http://standards.ieee.org/findstds/standard/8802-1AS-2014.html>
- [8] Hirschmann IEEE 1588 White Paper, available at <http://belden.picturepark.com/Website/?Action=viewPdf&AssetId=8050>
- [9] <http://www.ieee802.org/1/pages/802.1AS-rev.html>
- [10] <http://www.ieee802.org/1/pages/802.1ci.html>
- [11] <http://standards.ieee.org/findstds/standard/802.1AE-2006.html>



Belden Competence Center

As the complexity of communication and connectivity solutions has increased, so have the requirements for design, implementation and maintenance of these solutions. For users, acquiring and verifying the latest expert knowledge plays a decisive role in this. As a reliable partner for end-to-end solutions, Belden offers expert consulting, design, technical support, as well as technology and product training courses, from a single source: Belden Competence Center. In addition, we offer you the right qualification for every area of expertise through the world's first certification program for industrial networks. Up-to-date manufacturer's expertise, an international service network and access to external specialists guarantee you the best possible support for products. Irrespective of the technology you use, you can rely on our full support – from implementation to optimization of every aspect of daily operations.



About Belden

Belden Inc., a global leader in high quality, end-to-end signal transmission solutions, delivers a comprehensive product portfolio designed to meet the mission-critical network infrastructure needs of industrial, enterprise and broadcast markets. With innovative solutions targeted at reliable and secure transmission of rapidly growing amounts of data, audio and video needed for today's applications, Belden is at the center of the global transformation to a connected world. Founded in 1902, the company is headquartered in St. Louis, USA, and has manufacturing capabilities in North and South America, Europe and Asia.

For more information, visit us at:

www.belden.com

www.beldensolutions.com

follow us on Twitter [@BeldenIND](https://twitter.com/BeldenIND)