# TSN Influences on ODVA Technologies: IEEE-802.1, Avnu, IETF

Presented at the ODVA
2017 Industry Conference & 18th Annual Meeting
February 21-23, 2017
Palm Harbor, Florida, USA

Steve Zuponcic
Technology Manager
Rockwell Automation

Rudy Klecka
Technical Lead
Cisco

Mark Hantel
Project Engineer
Rockwell Automation

Paul Didier
Solution Architect
Cisco

## ABSTRACT
Time Sensitive Networks standards are being considered for adoption by ODVA into EtherNet/IP. This paper will provide a status and update to the ODVA community on the underlying TSN groundwork being done by other standards groups, and the potential impacts of this work on EtherNet/IP.

1) Review the current TSN standards status in IEEE,
2) Review the work in IETF DETNET to create wide-area deterministic networks
3) Review the work of Avnu developing interoperability and certification standards for TSN, including the Theory of Operations for TSN, a systems design document.
4) Review current ODVA SIG activities in response to new TSN technologies.

The paper will review the standard status and content that make up IEEE 802.1 TSN, especially:
• IEEE 802.1Qcc Stream Reservation Protocol (SRP) Enhancements and Performance Improvements;
• IEEE 802.1Qbv Enhancements for Scheduled Traffic;
• IEEE 802.1Qbu Frame Preemption;
• Precision Time (IEEE1588 and IEEE802.1AS)

This paper provides a summary for the following TSN standards:
• IEEE 802.1Qci Per-stream filtering and policing;
• IEEE 802.1CB Frame Replication and Elimination for Reliability;
• IEEE 802.1Qca Path Control and Reservation;

Finally, this paper discusses how EtherNet/IP and TSN standards might combine to provide a new generation of “standard, unmodified, Ethernet” with features and capabilities not previously achievable.

## KEYWORDS
ODVA, EtherNet/IP, CIP, TSN, Time Sensitive Networks, SDN, Software Defined Networking, IEEE 802.1, 802.1, Scheduled Ethernet, Industrial Ethernet, CNC, Centralized Network Configuration, CUC, Centralized User Configuration, Network Calculus

## TABLE OF CONTENTS



## DEFINITION OF TERMS

*CUC:* Centralized User Configuration engine. In an industrial system, the CUC is most likely associated with the PLC programming software for offline configuration, or, with the PLC during runtime. It is in these components where network information is already known.

*CNC:* Centralized Network Configuration policy engine. The controller that receives all network and policy requirements and calculates both the network calculus as well as the schedule for the network.

*Scheduled Traffic:* Traffic that is managed via the 802.1Qbv specification which defines a time-based shaper for bounding both latency and variations in data deliver.

*Rate-constrained traffic:* All traffic that uses traditional quality of service (QoS) methods of prioritization and that is not scheduled traffic.

*Best effort traffic:* Traffic that does not have any guarantees that data is delivered nor is given a QoS level or specific priority.

*Network Calculus:* The mathematics for calculating network loading and for modeling the network characteristics given payloads, maximum latencies per stream, and requested packet intervals.

*RPI:* Requested Packet Interval. This is the requested sampling frequency for data exchange across a connection.

# 1 INTRODUCTION

Time Sensitive Networking (TSN) refers to a set of standards that are being driven and developed by the IEEE 802.1 Time Sensitive Networking Task Group. This task group was originally called the Audio / Video Bridging Task Group, where the focus of moving audio and video information across local area networks without latency, phase shift, or packet loss was the primary emphasis. In 2012, the name was changed to reflect the wider push for these technologies and standards toward markets outside of the Audio/Video Bridging market. The group focuses on the technologies required to move time sensitive data over layer 2 networks in a deterministic fashion.

The development of Time Sensitive Networking (TSN) technologies has rippled across many industry segments and many markets. The initial thrust of these efforts was highly technical and aimed at solving the specific problem of guaranteed delivery of data across “standard Ethernet.” The extensions defined, however, create mechanisms that can potentially disrupt existing technologies and leave existing installations vulnerable to new traffic if not managed properly. Essentially, TSN technologies create new prioritization mechanisms that have the capability to override traffic that had previously been given highest priority under traditional QoS rules.

So how to introduce a new level of capability and higher levels of priority for new traffic without disturbing existing technologies and applications? This is the question that has consumed much of the market dialogue and which has been a key point of conversation in the standards communities. This white paper provides an overview on TSN technologies, the influencing bodies that are shaping its implementation and direction, and addresses key issues and concerns relative to existing ODVA technologies.

# 2 WHAT ARE THE TSN STANDARDS?

It's important to note that the TSN standards and technologies represent approximately 20 new projects that contribute to the IEEE-802.1 portfolio of standards. There is not simply one standard that is being launched but many that have impact on the way that traffic may be managed in the future. The following list these new standards at a high level:

IEEE 802.1 TSN – Completed

- IEEE 802.1AS-2011 – timing and synchronization
    - A profile of IEEE1588 that is tailored to IP Networks with high precision. It uses the peer-to-peer delay mechanism.
- IEEE 802.1Qat-2010 – stream reservation protocol
    - A distributed method for configuring stream reservations through a layer two bridged network. It mainly relies on MRP.
    - Included in 802.1Q-2012
- IEEE 802.1Qav-2009 – forwarding and queuing for time sensitive streams
    - A credit based shaper for bounding latency as well as variations in delivery.
    - Included in 802.1Q-2012
- IEEE 802.1BA-2011 – audio video bridging systems
    - A profile for use of Qav/Qat in Audio/Video Bridging systems.
- IEEE 802.1Qbu & 802.3br – frame preemption (work being done in P802.3br)
    - Allows high priority traffic to interrupt lower priority traffic so that it may be immediately transmitted without waiting for lower priority traffic to completely finish transmitting. Defines mechanisms that allow for lower priority traffic to be paused, fragmented, and reassembled after high priority traffic (aka preemptive traffic) has finished transmitting.
- P802.1Qbv – enhancements for scheduled traffic
    - A time-based shaper for bounding latency as well as variations in delivery.

IEEE 802.1 TSN – Almost Done (IEEE Sponsor Ballot)

- P802.1CB – frame replication and elimination for reliability (seamless redundancy)

    - Contains Stream Identification functions and methodologies to identify multiple paths through a network. Streams can then be duplicated for redundancy/reliability over guaranteed unique paths.
  - P802.1Qci – per stream filtering and policing
    - Improves the robustness of a Qbv (scheduled) network by policing and filtering traffic based on its configured resource requirements.
  - P802.1Qch – cyclic queuing and forwarding
    - Provides the algorithms to synchronize transmission of streams in a Qbv network through layer two bridges and endpoints to minimize latency and jitter.

IEEE 802.1 TSN – In Process

- P802.1AS-REV
  - Project targeted at providing enhancements to 802.1AS specification. A goal is to specify redundancy and fault tolerance for 802.1AS time-aware systems.
- P802.1Qca – path control and reservation
  - Allowing the use of distributed protocols to provide explicit path control, bandwidth and stream reservations.
- P802.1Qcc – stream reservation protocol enhancements
  - A centralized method for managing a network implementing TSN
- P802.1CM – Fronthaul Specification
  - A wireless carrier profile for TSN concentrating on QoS and preemption.

IEEE 802.1 TSN – Just Getting Started

- P802.1Qcr Asynchronous Traffic Shaping. [aka, UBS, or Urgency Based Scheduler]
- MRP++ (updated reservation protocol)
- There are additional IEEE standards that are less relevant

Without going into detail on each of the specifications given above, it's not difficult to see that TSN technologies represent numerous capabilities and services that will change the way that networking is accomplished moving forward.

## 2.1 IEEE 802.1Qbv: Scheduled Networking

Although there are numerous specifications that make up the TSN complement listed above, the majority of work expended by the various standards and compliance organizations at a system and interoperability level is being focused on the 802.1Qbv standard, or scheduled Ethernet. The reason this particular technology has such a draw is because it is believed that the scheduled Ethernet technology can provide the best mechanism for control of streams and the absolute knowledge of when data is sent and how it is being coordinated with other streams on the wire. This particular technology can provide a level of control that no previous, open standard, Ethernet technology has been able to provide. It allows for an absolute management of data to the extent that modeling and simulation of network traffic becomes practical, and the number of unknown parameters is reduced to solvable terms.

There is a challenge, however, to the introduction of scheduled Ethernet as a new, IEEE-based, standard Ethernet technology. By its very nature, scheduled Ethernet must become the highest priority form of traffic on the wire. When merging this technology into existing installations, migration mechanisms are required to allow for appropriate blending of “old” traffic with “new” traffic. From an ODVA perspective this is especially true of EtherNet/IP traffic because EtherNet/IP does not support a scheduling mechanism. Consider Figures 1 & 2 below:

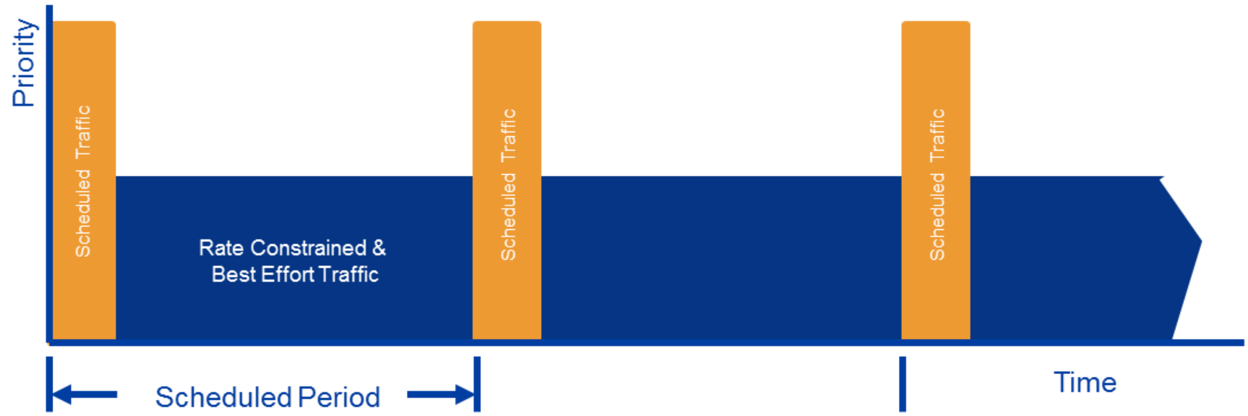


**Figure 1: Scheduled traffic with light duty cycle**

In figure 1, scheduled traffic is shown as the highest priority traffic on the Ethernet network. All previously existing, standard, prioritization mechanisms are defined as "rate constrained", while unmanaged traffic is defined as "best effort." In this model, today's EtherNet/IP traffic would fall into either the rate constrained or best effort traffic categories – depending upon the prioritization mechanisms that were used to configure the traffic during network commissioning. All previously defined prioritization mechanisms are defined as lower priority than scheduled traffic. The use case in figure 1 shows the coexistence of scheduled traffic with other traffic types - and here – the scheduled traffic has a relatively light duty cycle within the frame of the schedule itself.

The use case presented in figure 2, however, is different. Here, the duty cycle, or percent loading on the wire is fairly high for the scheduled traffic, compared to that of the rate constrained or best effort traffic. The question, however, is: Is this problematic? It really depends on the application and the requirements of the system.

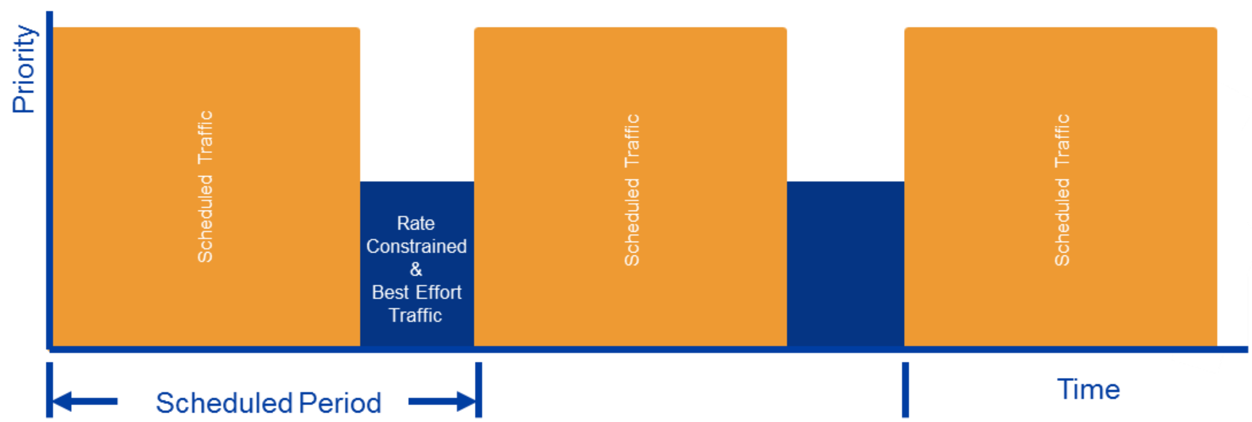


**Figure 2: Scheduled traffic with heavy duty cycle**

The coexistence of scheduled Ethernet traffic with non-scheduled Ethernet traffic can be managed with the proper understanding of all traffic requirements at a system level. That is, if all paths are defined, payloads are known, RPI's are determined, and maximum data delivery latencies are provided, then the aggregation of all streams can be calculated to determine if successful operation of the network will be accomplished. If calculations determine that planned data streams cannot be accommodated – then the user can use this information to modify the network by either adding additional paths in the topology, by increasing wire speed, or, by changing RPI's or loading in the system.

# 3 THE BIGGER PICTURE: SYSTEM BASED NETWORK CONFIGURATION AND COMMISSIONING

A core precept of the TSN value proposition is that all network communications are managed so that there is a guarantee for performance and for data delivery. To accomplish this, all devices need to participate in traffic planning by publishing to, or notifying, a Centralized Network Configuration policy engine (CNC) of their traffic requirements for the connections involved. The CNC, which is a new system level function introduced by the TSN standards organizations, will calculate the best possible solution to accommodate all the traffic flows between all connections in the network. It has the role of configuring the infrastructure components in the network (i.e., the bridges) to accommodate the traffic flows, given some combination of the services and standards listed above. If the CNC is not able to solve for a configuration that can meet the performance and loads of all the traffic streams in the subnet, it will notify the user so that topology, performance requirements, or loading might be modified. If any set of devices do not participate in this planning process, the risk is that the devices that do not participate may interfere with planned traffic, or, more likely, that they will not have their own traffic streams properly serviced for successful operation in the application.

This aspect of the TSN solution dramatically changes the workflow for designing and planning networks. Traditional methods for network commissioning have depended on component by component configuration with no real focus on a holistic solution. One of the TSN value statements is that network calculus and network planning become a part of the solution toward managing traffic and guaranteeing performance. In this new paradigm, payload, sampling frequency, and maximum latency are all communicated to the Central Network Controller (CNC) so that there is a single place to calculate flows and to configure bridges and infrastructure to meet these demands. This new capability is powerful in that it can enable offline simulation of the network even before hardware is procured or cabling is routed. If designs prove inadequate or a solution is not achievable given system constraints, then network design can be modified to accommodate the requirements of the system.

Consider the following high level system diagram which shows the roles and functions of key components in the TSN architecture:

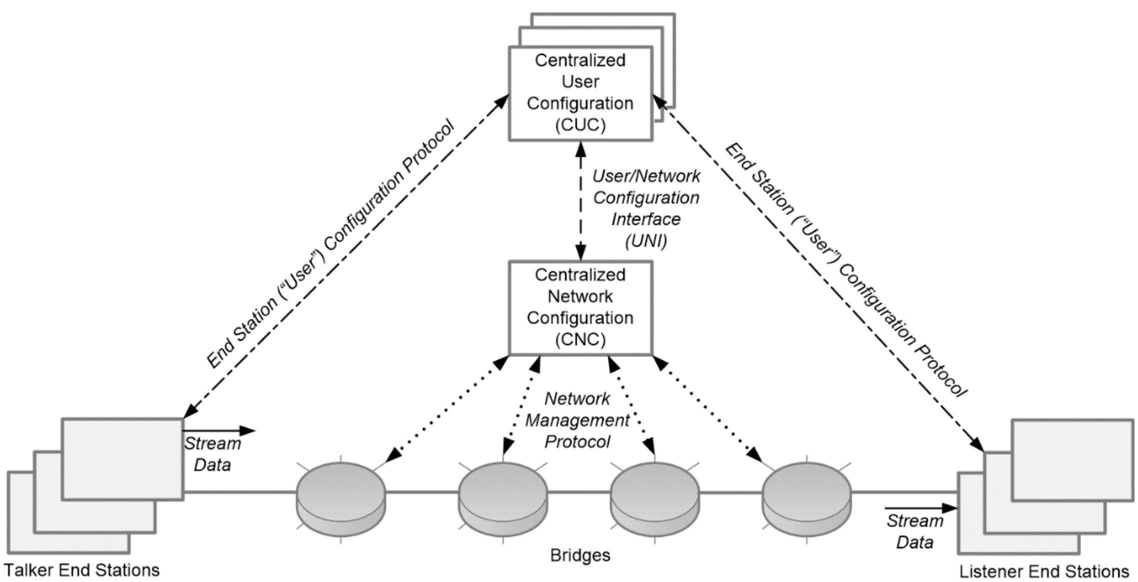


**Figure 3: Centralized Network Configuration**

In this diagram, the CNC function has the role of configuring all bridges, or switches, in the infrastructure. It communicates with the Centralized User Configuration tools (CUC) to receive the information about the traffic flows that need to be serviced in the network. It is this link between the CUC and the CNC where the information about payload size, RPI's and maximum latency are communicated to reflect the needs of the application. The communications between the CUC and the CNC is referred to as the User/Network Configuration Interface or UNI. The attributes that are shared across this link for proper stream management are defined in the 802.1 Qcc specification. In an industrial system, the CUC is most likely associated with the PLC programming software for offline configuration, or, with the PLC during runtime. It is in these components where network information is already known; it is in these places where I/O connections, payloads, and system latency have already been defined in order to achieve the needs of the industrial application.

## 3.1 POLICY BASED PRIORITIZATION

The introduction of system level configuration tools for design and commissioning carries the prospect for true control and management of the network as well as the application. As in all networks, the policies that are established by the network administrator are always a function of the needs of the application and the requirements of the system. Policies are established to provide prioritization to the traffic that is most important in the application, and to ensure that the most important traffic is given precedence over competing streams in the system. And not all systems are designed the same; they won't all give the same level of priority to the same traffic types. For example, a telephone system is likely to give highest priority to VOIP while an industrial control system may give highest priority to its motion control traffic.

This element of policy based management gives complete control to any person managing the network or the system. Moreover, if this management is enabled through system level software that interfaces directly to the CNC, the network administrator is given an extremely powerful tool for interacting with the system. As previously discussed, the CNC is given all requirements for the predefined streams given payload sizes, RPI's and maximum latency. If policy information is also provided as an input to the CNC, then the system has all the information it needs to ensure that the highest priority traffic is given access to the wire. *This is regardless of the mechanism that is used to place data on the network! If there is scheduled traffic and unscheduled traffic – unscheduled traffic may still be given highest priority over scheduled traffic simply because the policy dictates it!* As an example, ODVA's CIP Motion traffic is not scheduled traffic from an 802.1 Qbv perspective; however, it may share the same network as traffic from a video camera which may be scheduled. Regardless of the network access methods used, if the user policy dictates that the CIP Motion traffic is more important than the video traffic, the CNC and corresponding system software must calculate and allocate networking resources to allow the CIP Motion traffic access to the wire. If the CNC cannot successfully solve for this outcome given all constraints – an error is returned to the network administrator so that reengineering of the network or the network parameters can be accomplished.

Policy-based prioritization infers that application specific priorities are provided to the system via user interaction with the CNC.

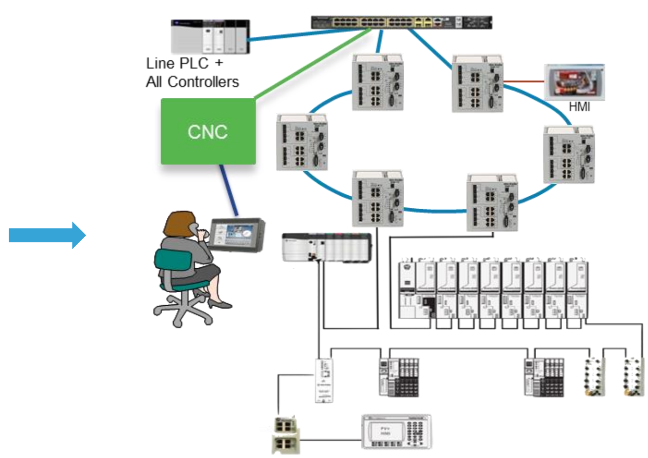


**Figure 4: Policy-based prioritization and system configuration**

Policy-based prioritization, then, becomes the counter balance for all network access methods that are defined in the TSN standards – or in any previously defined Ethernet standards for that matter.  By coupling the information of the policy with the tools used to provide access to the network, true "fairness on the wire" can be accomplished to allow for a balanced convergence of traffic on any Ethernet network.

### 3.2 CNC FOR SCHEDULED AND UNSCHEDULED TRAFFIC CONFIGURATION

The Centralized Network Configuration policy engine has the responsibility to configure the network resources based on the policies given as well as the capabilities and constraints of the infrastructure it is configuring.   There are two broad calculating engines that are potentially contained in the CNC:  One is a scheduling engine (SE) for solving the scheduled part of the network configuration policy.   The other engine is a network calculus engine (NCE) which is used to solve the rate constrained (traditional QoS) part of the network configuration policy.

The CNC policy engine would normally contain both the SE as well as the NCE to find the best solution for all the parameters given to it.  It's important to note that the CNC policy engine can manage either scheduled traffic, or, rate constrained traffic, or both, to create a successfully managed network.

## 4 AVNU ALLIANCE

The Avnu Alliance organization is the certifying body that develops interoperability specifications for the TSN standards.  The Avnu organization consists of a group of companies from various industries and market segments.  These are the professional audio/video market, the automotive market, the consumer market, and the industrial market.  Having begun in the audio/video market, the first Avnu conformance specifications have been written to certify audio/video products, and there are many products that have been certified and that have received the Avnu certification mark.

The industrial market segment is relatively new to the Avnu Alliance. The organization has begun its work in developing certification standards for the industrial control sector.  One of the challenges in applying TSN technologies in the industrial sector is that there are multiple existing technologies and standards which are long established and have large installed bases.

As an example, ODVA today has no notion of scheduled traffic on Ethernet (reference TSN specification 802.1Qbv) and it utilizes standards different than the Avnu specified 802.1AS mechanism for clock synchronization.  A key discussion area within Avnu is to define TSN-based industrial profiles and conformance specifications which allow for the coexistence of predefined technologies and for the protection of existing installations. There are many members of ODVA that are also members of Avnu that are working toward this goal.  This work is no small task; it requires diligent and exacting attention to the mapping and merging of these technologies for the sake of migration and network convergence.

### 4.1 INDUSTRIAL THEORY OF OPERATION DOCUMENT

Today, the Avnu Alliance has released its first revision of the "Theory of Operation" document which describes the baseline, foundational, technologies that form today's version of TSN.  This document will continue to be modified and augmented to include those elements that provide the ODVA based technologies a migration path forward.  As this document evolves, it will become the document that helps to define the industrial conformance standards going forward.

## 5 INDUSTRIAL CONTROL OVERVIEW

Modern industrial networks combine the disciplines of both information technology (IT) and operational technology (OT) to meet the requirements of industrial applications   The applications served by the industrial control sector are sophisticated processes requiring complex and modularized designs that often require modification and augmentation during runtime.  These applications require high performance, determinism, and predictability; their infrastructures demand highly robust and reliable designs both in hardware durability and in resiliency features to support high-availability needs.

The networking legacy in the industrial control market is long established with roots that extend to multiple industrial technology standards as well as to multiple vendor-specific, proprietary technologies. The industrial installed base that has settled on these technologies is a conservative community of manufacturers that move carefully and slowly from one generation of technology to the next. Solutions need to be proven and ROI is carefully calculated before investments are made. The move from any current position is always done through migration and evolution in order to protect installed assets and to maximize profitability for any given change in architecture.

Prior to the advent of TSN technologies, there hasn't been much motivation to drive these differing industrial solutions to a more unified approach. The TSN conversation, however, has driven an awareness that future networking solutions will require a holistic approach and a true system perspective. The future network requires a comprehensive answer that includes all elements - infrastructure and end stations alike – to be included in the final solution. Moreover, in order to solve the complete Industrial Internet of Things (IIoT) problem statement, this comprehensive solution will require the collaboration of the existing, differing, technologies and standards to morph toward solutions that can allow their inclusion in this new eco-system. If any specific technology, vendor, or product does not participate in this new paradigm, there is no way for "the system" to accommodate or plan for the traffic that needs to be managed to or from these products. The result is that these products will not be properly served in the overall design.

# 6 Industrial Use Case

Consider the following representative architecture of an industrial control application and associated network design:

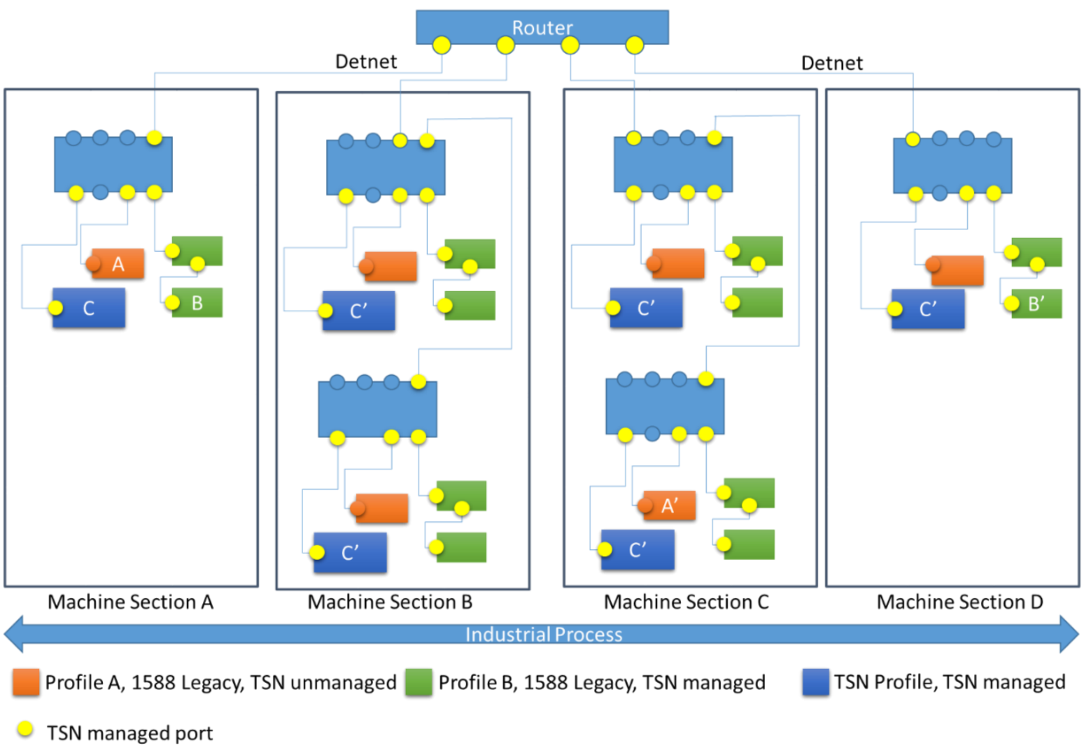


**Figure 5: The Industrial Use Case with Legacy Technologies and Devices**

In this system, an industrial process is controlled by a single machine which consists of four different sections of machinery. Each section of machinery is delivered to the process via different OEMs, each of which have different specialties in the parts of the process that they control. The end user has 7 manufacturing sites around the globe. There are 15 of these machines per site, and each machine IP addressing scheme is identical to the other machines in the same manufacturing facility, which matches the other addressing schemes of the other facilities. Moreover, each machine section is a subnet with a unique VLAN so that the equipment can be constructed modularly and functional organization is achieved. All sections of the machine are synchronized and coordinated to produce final product and relevant events are timestamped so that data on the manufacturing floor can be correlated against data in the MES system and data from the supply chain. The entire manufacturing facility uses the same understanding of absolute time and all events are related back to the common notion of "wall clock time." Additionally, this use case acknowledges the pre-existence of products and technologies that utilize other 1588 based solutions. In this case, component A in Machine Section A may communicate with component A' in Machine Section C; Component B in Section A may communicate with component B' in Machine Section D. Finally, component C in Machine Section A produces data to be consumed by C' components in sections B, C, and D. The implementation of time bridges or time gateway mechanisms via the switches provides a migration path for an extended community of existing products and technologies to be included in the wider TSN value proposition. In this use case, components A and A', and B and B' would require time gateway translations.

## 6.1 THE TSN VALUE PROPOSITION, MIGRATION MECHANISMS, AND PATH FORWARD

TSN technologies provide a path toward a system-wide view of the architecture. They allow, for the first time, standardized methods for network-wide configuration and commissioning. TSN technologies create the avenue for offline simulation and modeling of the network and the architecture. However, TSN technologies also demand the convergence of existing Ethernet technologies toward a unified set of standards and techniques.

In contrast, the use case in figure 5 is a very real depiction of today's industrial situation. There are a wide variety of existing standards and technologies that make up the industrial landscape. More importantly, there is a very large base of installed equipment which utilizes these existing, differing, technologies and which are producing products for numerous manufacturing companies globally. It is impossible for these installations to just "jump" into the TSN eco-system, or for industrial control suppliers to switch, bumplessly, toward these new solutions.

This suggests, then, a motivating action for both the driving organizations of TSN-based technologies as well as the adopting organizations of TSN-based technologies. In order to move forward, it is imperative that migration paths and mechanisms are defined which allow for the merging of existing technologies toward TSN-based systems. Without these migration mechanisms, existing installations and products have no method to participate and are left technologically marooned. This hurts the advancement of TSN-based technology because adoption is stunted; it also hurts the adopting community's progress because it blocks a pathway which provides for features that the market has been demanding for many years.

There are several examples of migrating mechanisms that have been discussed in the various standards organizations. One such mechanism is a "time bridge." A time bridge, sometimes referred to as a time gateway, is a mechanism for translating time across different time domains. Today, the Avnu Alliance has standardized on the 802.1AS profile for time synchronization in TSN systems. There are many other time profiles in existence, including the IEEE-1588 default profile utilized by ODVA's CIP Sync time object. A time bridge that synchronizes time between these two domains would allow for existing systems and devices to participate in the broader TSN-based architecture without making hardware changes.

Other migration mechanisms might include "TSN adapters" or switches that can blend best effort and rate-constrained traffic with scheduled traffic.

# 7 SCOPE AND TIMING

It's important to note that the industrial use case given in figure 5, above, is typical for a very wide range of industrial applications. This use case illustrates the need for solutions at layer 2 (switching), layer 3 (routing), and for time bridging functions.

Today, the IEEE-802.1 TSN Workgroup focuses on layer 2 in the architecture and this is where the majority of the TSN work is being defined. Here, there is work complete enough to begin implementation in a system or in a product. Additionally, Avnu has defined compliance specifications sufficient enough to certify some products for this level of operation. Avnu's "Theory of Operation" document focuses on the technology in this part of the architecture and will expand its scope in future versions to include other areas as those technologies become defined.

Time bridging, as discussed previously, is shown in the layer 2 switches and provides the mechanism for migrating existing time domains and PTP profiles into a TSN system. The time gateway function has not been identified as a required work item for any standards communities as yet. Individual suppliers could develop these bridge functions as solutions for the market.

# 8 The Internet Engineering Task Force (IETF)

Layer 3 (routing) functionality is currently being designed by the Internet Engineering Task Force (IETF) and it essentially extends the capabilities of the 802.1 TSN specifications into layer 3. This work, often called "Detnet", (deterministic network) is not yet complete, at the time of this writing.

# 9 ODVA's Distributed Motion and Time Synchronization SIG

In November of 2015, the ODVA Distributed Motion SIG changed its name to the Distributed Motion and Time Synchronization SIG (DM & TS SIG) in order to respond to the market's movement toward Time Sensitive Networking Technologies. In this transition, the SIG extended its charter to include a focus on the integration of Time Sensitive Networking technologies into the ODVA standards.

A primary focus and priority for the DM & TS SIG is to develop approaches and migration mechanisms that minimize impact of TSN integration on current product designs and existing installations. In order to facilitate this goal, the SIG was asked to focus on the following areas as a starting point:

- Frame Preemption (IEEE 802.1Qbu & 802.3br)
- Gigabit Ethernet
- Stream Reservation Protocols (IEEE.802.1 Qcc) and related technologies

These particular technologies were targeted because early analysis revealed that the combination of frame preemption combined with Gigabit Ethernet contributed dramatically toward system performance, determinism, and throughput as compared to scheduling alone. This analysis proved that high performance applications like motion control could benefit greatly from Gigabit Ethernet for bandwidth – and from preemption for reduction of latency behind large packets.

## 9.1 Application Analysis

At a high level, the analysis which was conducted consisted of a CIP Motion application with a single controller and multiple axes connected through a switch which acted as a distribution mechanism for multiple DLR segments that fanned out to 25 axes per segment. The system was analyzed at both 100 Mbit/s and 1 Gbit/s wire speeds. In addition, frame preemption and scheduling models were introduced to see the comparative differences in the number of axes that could be supported as each of these variables were modified in the architecture.

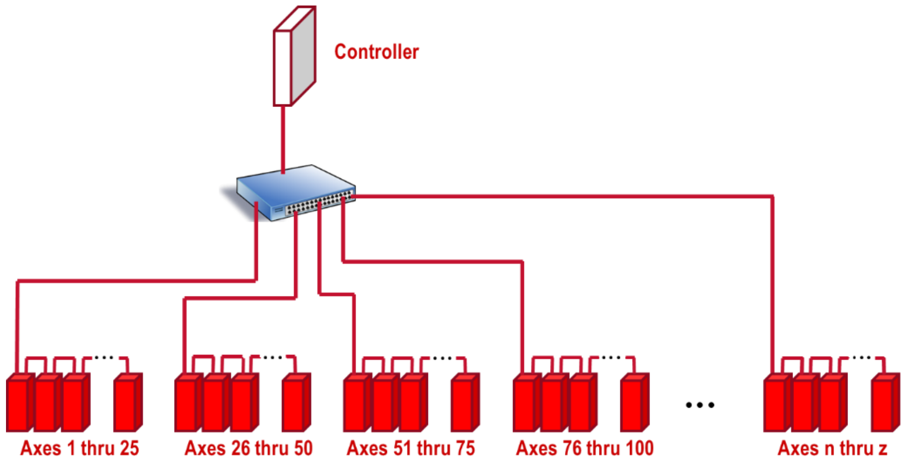


**Figure 6: Distributed Linear Motion Architecture**

The results, from a raw performance perspective, measured in total number of axes that can be supported, are given in the following table:

| | No TSN Functions | Preemption Enabled | Scheduling Enabled |
|---|---|---|---|
| **100 MBit/s** | 17 Axes | 35 Axes | 37 Axes |
| **1 GBit/s** | 411 Axes | 430 Axes | 432 Axes |

**Table 1: Supportable Axes at Different Wire Speeds with & without TSN features**

This analysis simply shows that, from a raw performance perspective, the effect of wire speed far outweighs the contributions that can be gained from either preemption or scheduling alone. However, preemption and scheduling both have the benefit of managing traffic when large packets are introduced to the network that are lower priority than CIP Motion traffic, (or any other traffic that is determined to be high priority.) Using preemption and scheduling traffic both dramatically reduce the latency of large packets on the wire. The following table shows the impact of a large 1518 byte packet introduced on the network as a function of wire speed, preemption, and scheduling. (Note, that when preemption is enabled, a maximum of 124 bytes of the lower priority frame will go through the buffer before the higher priority packet is passed through.)

| | 100 Mbit/s | 1 Gbit/s |
|---|---|---|
| **1518 bytes (w/o preemption)** | 123 microseconds | 12.3 microseconds |
| **124 bytes (w/ preemption)** | ~10 microseconds | 1.0 microseconds |
| **0 bytes (scheduled traffic)** | 0.0 microseconds | 0.0 microseconds |

**Table 2: Latency behind large, lower priority, packets with and without preemption or scheduling**

## 9.2 Coexistence with Scheduled Traffic & Integration of 802.1 Qcc

Despite the performance benefits shown in this analysis, the market is slowly moving toward adoption of the IEEE-802.1 Qbv (scheduled Ethernet) technologies for several reasons. The ability to control streams and to model and simulate entire networks is simplified if scheduled times are provided for critical traffic. Like trains in a railway system, the tight scheduling of multiple trains across the same tracks to and from various destinations can be controlled for optimized throughput while avoiding collisions. Similar to allowing passengers to plan their journeys around known starting and ending times – the network simulation tools can now provide predictable models that allow network designers to know well ahead of commissioning whether their system will be successful or not.

This move by the market toward IEEE-802.1 Qbv means that existing products will need to participate in this new environment by publishing their networking requirements to the CNC and allowing it to manage and setup the infrastructure for successful management of all traffic. If any device proposes to share the same wire with other TSN-enabled devices, publication of requirements to the CNC is essential to prevent being "run over" by other traffic streams. This means that regardless of any decision to adopt or not to adopt scheduling as a part of the CIP architecture, new services to allow for the publication of stream requirements is a baseline "must" for coexistence with scheduled traffic.

The 802.1 Qcc standard is the stream reservation protocol that needs to be adopted within ODVA to allow products to publish their requirements to the CNC. This publication of information would be provided to the CNC by the CUC. In an industrial application, the CUC is likely to be either the application programming software for the controller in an offline situation, or, the PLC in an online situation. These components already have the stream requirements defined.

As EtherNet/IP products move toward the TSN domain, the inclusion of 802.1Qcc services becomes a baseline requirement for coexistence and the DM & TS SIG will begin definition of these services within the ODVA specification as these standards solidify within IEEE.

## 9.3 Current Status

The ODVA Distributed Motion and Time Synchronization SIG has identified those TSN-standards that need to be integrated into ODVA specifications, and also, it has identified those technology areas that require further development in order to facilitate the convergence of ODVA technologies with TSN technologies in the future. The following list details those areas:

- Continued development and promotion of Gbit Ethernet
- Inclusion of IEEE 802.1Qbu & 802.3br for frame preemption
- Integration of IEEE 802.1 Qcc and other stream reservation services for publication of network requirements to the CNC
- Definition, development, and promotion of time gateways to allow for bridging of various time domains and time profiles (E.g., 802.1AS to CIP Sync's IEEE-1588 default profile)

## 9.4 EtherNet/IP Uniquely Positioned for CNC integration using traditional QoS

It should be noted that EtherNet/IP is uniquely positioned to take advantage of the Network Calculus engine portion of the CNC, as discussed in section 3.2 above. Because EtherNet/IP has always been designed on a "standard, unmodified, Ethernet" platform, its use of standardized prioritization mechanisms puts it in a ready position for communicating requirements to the CNC long before other technologies are ready with scheduled solutions, or with the development of products that can manage the scheduling engine portion of the CNC. This means that ODVA members may benefit from early TSN participation by providing solutions that are already in the ODVA "sweet spot" of capabilities.