A User's Perspective on Wired EtherNet/IP Network Architectures

Gary Workman
Consultant

John Rinaldi
Director of Creating WOW!
Real Time Automation

Presented at the ODVA
2025 Industry Conference & 23rd Annual Meeting
March 19, 2025
Clearwater Beach, Florida, USA

Abstract
From a user's perspective, an EtherNet/IP network is a control system network.

An end user with that perspective could certainly inquire of an EtherNet/IP authority, "How should I architect an EtherNet/IP control system network?" But asking that question itself raises other questions that are worth exploring. Is an EtherNet/IP network truly a control system network? Is the "user's perspective" qualifier necessary? Why would users think an EtherNet/IP network is a control system network?

All these interesting questions are rooted in the two fundamental questions examined at the beginning of this paper:

What is an EtherNet/IP network? What is a control system network?

Keywords
EtherNet/IP, EtherNet/IP Networks, Control System Networks, ControlNet, Information System Networks, EtherNet/IP Network Architecture, Control system design, EtherNet/IP DLR

What is an EtherNet/IP Network?
EtherNet/IP is a trademark of the ODVA organization, the organization which owns and publishes the EtherNet/IP specifications. Reviewing those specifications would appear to be a good place to begin searching for the definition of what constitutes an EtherNet/IP network. Pursuing that path leads to discovering that neither the EtherNet/IP specifications nor the associated ODVA library Publication 35 – "EtherNet/IP Network Infrastructure Guide" directly specify or define what constitutes an EtherNet/IP network. Rather, the EtherNet/IP specifications specify an EtherNet/IP conformant device – a device using the CIP application protocol serviced by a TCP/IP protocol stack to communicate implicit and explicit messages via an Ethernet network interface.

While the EtherNet/IP specifications don't directly describe what constitutes an EtherNet/IP network, perhaps the specification of what an EtherNet/IP device is can be used to infer what an EtherNet/IP network is.

The original EtherNet/IP specification can be traced back to the ControlNet specifications. The direct predecessor to the EtherNet/IP specification text was text that first appeared in an appendix of the ControlNet specification. That appendix described how CIP explicit messages and other TCP/IP application protocols could be relayed to devices on or across an Ethernet information network. Using a TCP connection via an IP router and the unscheduled bandwidth of a ControlNet network, explicit messages could be exchanged with any devices on an Ethernet network that implemented a CIP explicit messaging capability. The initial EtherNet/IP specification was published by the ControlNet International organization after the concept of an implicit message connection for exchanging control signal traffic between devices on an Ethernet network was developed. Even though the EtherNet/IP specifications weren't published until after including text describing an implicit messaging connection for communicating control signal data, support for the implicit messaging capability was specified as an optional feature of an EtherNet/IP conformant device.

Consider the first three figures shown below. Figure 1 depicts two devices on separate ControlNet networks exchanging data via an explicit messaging connection across a shared media Ethernet information system network. It obviously does not depict an EtherNet/IP network because no EtherNet/IP capable device is shown in the figure.

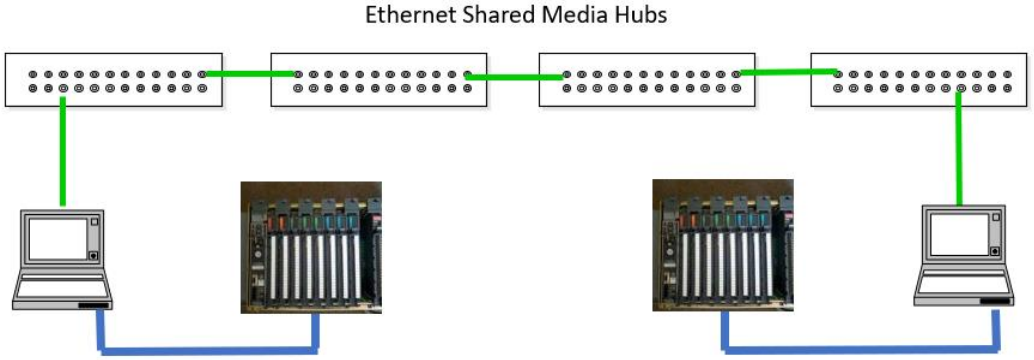

Figure 1 - PLC-5s on separate ControlNets exchanging CIP explicit messages using ControlNet to Ethernet routers

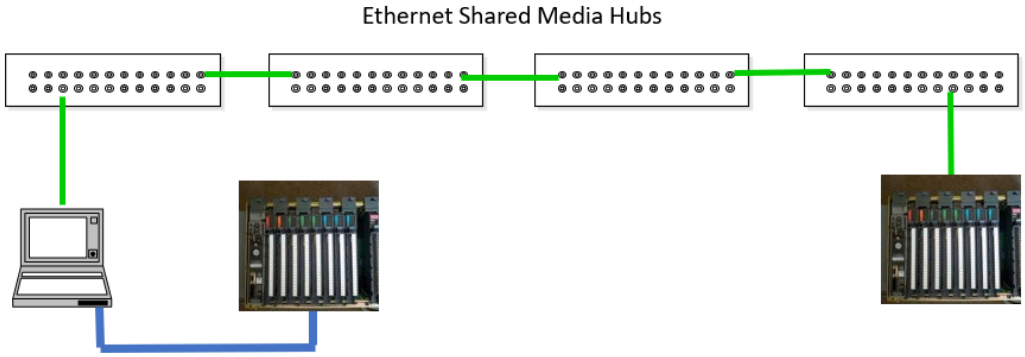

Figure 2 - PLC-5s on a ControlNet exchanging CIP explicit messages with a PLC on an Ethernet network using a ControlNet to Ethernet router

Figure 2 depicts a device on a ControlNet network exchanging data via an explicit messaging connection with a single device on a shared media Ethernet network. Assuming it takes at least two EtherNet/IP capable devices to create an EtherNet/IP network, this figure also does not depict an EtherNet/IP network. *[Interestingly,* Figure 2 *depicts the scenario described in the ControlNet specifications appendix text that formed the foundation for the EtherNet/IP specifications.]*

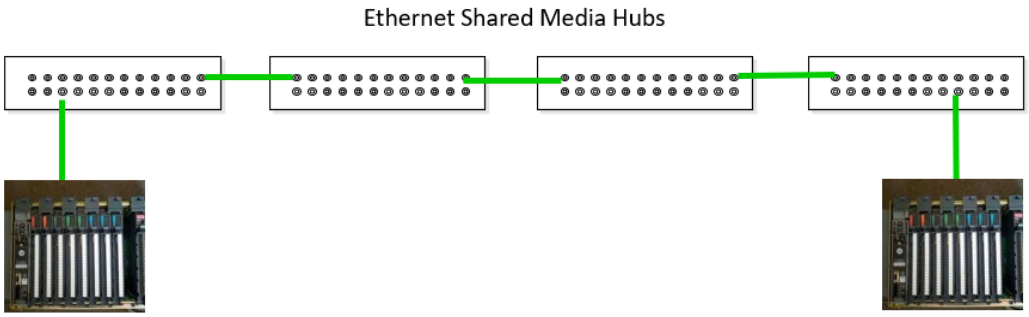

Figure 3 - PLC-5s in the same TCP/IP subnet exchanging CIP explicit messages via a shared Ethernet network

Figure 3 depicts two EtherNet/IP capable devices on the same shared media Ethernet network shown in the prior two figures. In Figure 3, the EtherNet/IP capable devices are both only capable of exchanging CIP explicit messaging traffic and are members of the same IP subnet. Assuming that the two devices communicate CIP traffic with each other, it would be hard to refute the conclusion that Figure 3 depicts an EtherNet/IP network.

Having concluded that Figure 3 depicts an EtherNet/IP network, it becomes both appropriate and reasonable to ask: Which portions of Figure 3 constitute the EtherNet/IP network? Are the Ethernet hubs part of the EtherNet/IP network? Are any other devices attached to the Ethernet hubs part of the EtherNet/IP network? If so, which devices? Is it only an EtherNet/IP network when the devices are actually communicating CIP traffic with each other, or is simply having the potential to communicate CIP traffic with each other enough to declare it to be an EtherNet/IP network?

Let's change a few of the conditions in Figure 3 and attempt to determine if we still have an EtherNet/IP network. What if the two EtherNet/IP capable devices shown in Figure 3 do not communicate with each other? What if they both separately communicate with the ControlNet capable device on the ControlNet network shown in Figure 2?

What if the two EtherNet/IP capable devices shown in Figure 3 do not belong to the same IP subnet? They both belong to the same shared media Ethernet network. Is that sufficient to have that Ethernet network or a portion of it be recognized as an EtherNet/IP network?

The answers to such questions require a further definition of what constitutes an EtherNet/IP network. The idea of exactly what constitutes an Ethernet network, let alone an EtherNet/IP network, can be a nebulous concept. The defining characteristics of an Ethernet network have changed as the Ethernet standards and technologies have evolved. The extent of a collision domain established the boundaries of an original shared media Ethernet network. A fully switched Ethernet network may not even possess a collision domain if all the links are operating in a full duplex fashion. An internet is a network of networks. A single Ethernet network can host multiple IP subnetworks and multiple VLANs (Virtual Local Area Networks).

To our awareness, the most network-related requirement in the EtherNet/IP specifications is a requirement that all EtherNet/IP capable nodes on an EtherNet/IP network must respond to the List Identity command broadcast on that network. Responses to the issuer of that command on an Ethernet network without an IP router can only come from EtherNet/IP capable devices belonging to the same IP subnet as the command issuer. Because of that behavior, this requirement appears to constrain the scope of an EtherNet/IP network to devices on an Ethernet network that are members of the same IP subnetwork.

Consider Figure 4. It depicts a single Ethernet unmanaged Ethernet switch with five EtherNet/IP capable devices connected to it. The isolated Ethernet switch is a single Ethernet network. Every device connected to it is a member of that Ethernet network. Any message broadcast by any device connected

to that switch will be sent to every other device connected to that switch. The devices belong to two different IP subnets. Are there one or two EtherNet/IP networks depicted in the figure?

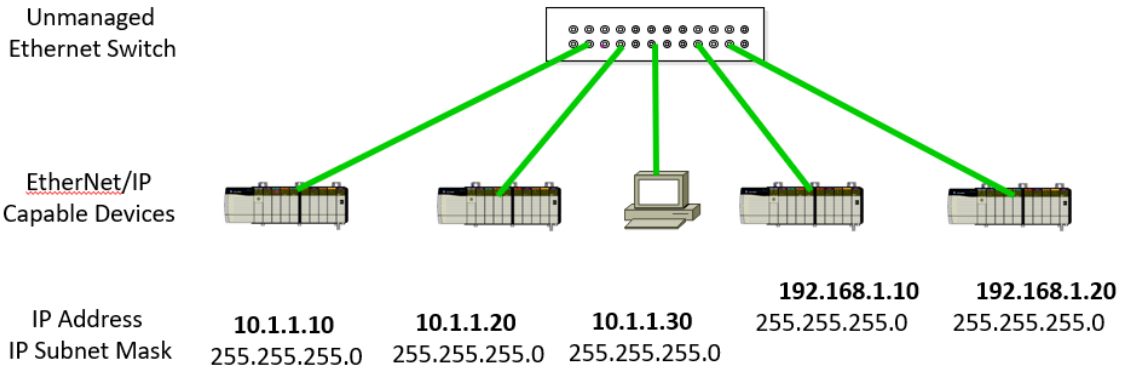

Figure 4 - Five EtherNet/IP capable devices connected to the same Ethernet network, each configured to belong to one of two TCP/IP subnets. How many EtherNet/IP Networks are depicted in the figure?

If responding to a broadcast EtherNet/IP List Identity command is used to define what constitutes an EtherNet/IP network, then Figure 4Figure 4 depicts two EtherNet/IP networks. The question of which network, if either, contains the unmanaged Ethernet switch remains unanswered.

The requirement that all EtherNet/IP capable nodes on an EtherNet/IP network must respond rather than a requirement that all nodes on an EtherNet/IP network must respond to a broadcast List Identity command implies that there can be non-EtherNet/IP capable nodes on an EtherNet/IP network. Without further clarification from an EtherNet/IP authority, it seems that all TCP/IP capable devices that belong to the same IP subnet as at least two EtherNet/IP capable devices could be considered EtherNet/IP network nodes.

What is a Control System Network?
Broadly speaking, for the purposes of this paper, there are two distinct categories of computer system networks – control system networks and information system networks. A control system network is fundamentally different than an information system network, even when they share a common networking technology foundation.

Control signal data exchange is the primary purpose of a control system network.
Information sharing is an afterthought of a control system network. Information sharing only occurs on a control system network if and when that network has the luxury of communication bandwidth in excess of the control system control signal data communication needs. If information sharing on a control system network interferes with the primary purpose of that network, it is a poorly designed control system network and should not be designated or recognized as such.

Information sharing is the primary purpose of an information system network.
Control signal data exchange on an information system network should only happen if and when the information system network meets sufficient performance, reliability, and repairability requirements to satisfy the control signal data exchange needs. Control signal data exchange on an information system network should only happen if and when that information system also dedicates sufficient bandwidth to ensure that it can satisfy the control signal data exchange timing performance needs. Those behaviors require an information system network to both recognize control signal data traffic and acknowledge that the control signal data traffic is among the most important information being communicated. Calling an information system network “help desk” to report a production stoppage issue is an unacceptable response mechanism for a control signal data exchange disruption on an information system network.
Control system networks are integral and vital components of machine control systems. A control system network exists and operates where the equipment connected to the control system network resides. A control system network is integral and indistinguishable from the production system and the devices

implementing it. The people responsible for maintaining and operating the production system are responsible for the maintenance and operation of the control system network and all its associated devices. A control system network is a control system network regardless of whether the network is connected to an information system network.

Early Investigations Proved Ethernet's Viability as a Control System Network
In the mid-1980s, as part of General Motors' Manufacturing Automation Protocol initiative, General Motors (GM) warned users in other companies that Ethernet's collision and collision recovery behavior made it an unacceptable foundation for control system networks.

Ethernet standards evolved. By the mid-1990s, full-duplex switched Ethernet made it possible to eliminate what used to be Ethernet's defining behavior – communication collisions. However, one key feature of those evolving Ethernet standards was the requirement for backward compatibility. Ethernet switches and end devices all had to support the half-duplex Ethernet operation and the collision recovery behaviors that cause undesirable performance characteristics for a control system network.

In 1997, GM re-investigated the idea of using Ethernet as the foundation for a controller-level control system network.

The key finding of that investigation is shown in

Figure 5.

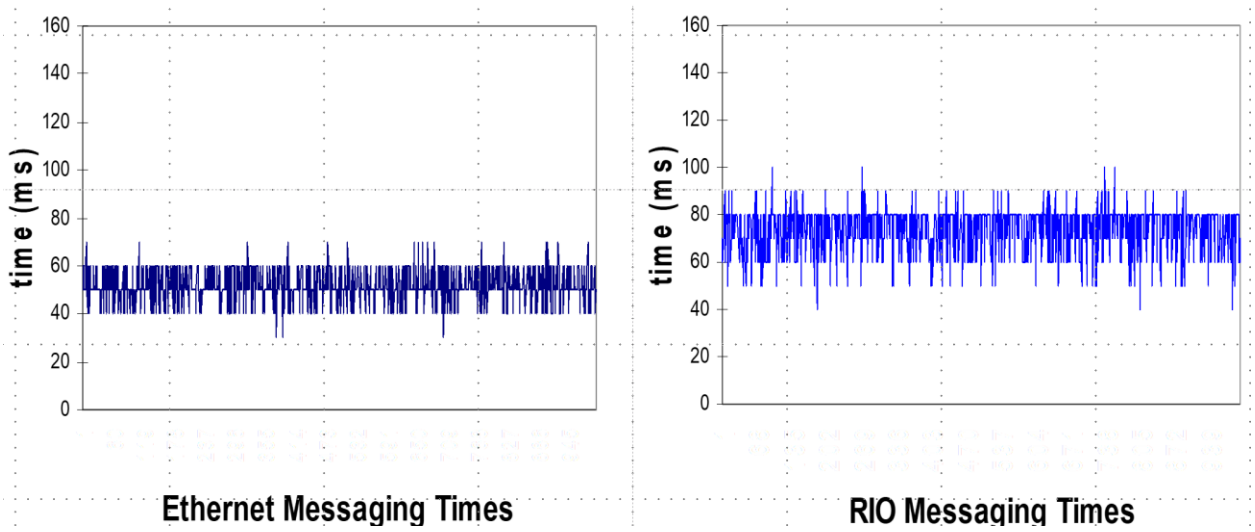

Figure 5 - 10-node network round-trip message delivery times

Figure 5 shows the round-trip message delivery times for exchanging one byte of control signal data between two PLCs on 10-node Ethernet and 10-node Remote I/O (RIO) networks. RIO was the deterministic, token-passing, control system network technology used by General Motors to build and assemble vehicles at that time. Even under unreasonably biased but plausible conditions, Ethernet-based message exchanges between PLCs on an isolated 10 Mbps Ethernet network outperformed an optimized Remote I/O (RIO) control system network connecting those same PLCs and exchanging the same control signal data. The tests were repeated with other models of PLCs with similar results.

Similar tests were run using personal computer-based controllers and those tests exposed performance-related issues associated with their Ethernet communication interfaces. In one of those tests, while the control programs were run under a real-time operating system, the Ethernet communication interface was

run under a Windows operating system, executing as a low-priority task of the real-time operating system. The Ethernet network communication of control signal data in those tests experienced occasional delays of a second or more due to that configuration of the Ethernet communication interface.

Another Ethernet interface performance issue was exposed when an Ethernet interface card with expanded memory advertised to enhance communication performance was used. In tests with that card, the control signal communication performance experienced frequent delays of several hundred milliseconds when simultaneously transferring a file. Apparently, several hundred frames of file transfer data were queued in the expanded memory and had to be transmitted before the next control signal data message could be sent.

These investigations concluded that Ethernet – if properly implemented – could effectively communicate control system traffic.

The investigation concluded that "A proper implementation" of an Ethernet-based control system network requires 1) that a device's Ethernet network interface must be able to identify and preferentially handle control system traffic and 2) a proper selection and configuration of network infrastructure components - industrialized cables and industrialized Ethernet switches operating in full-duplex fashion.

Is an EtherNet/IP Network a Control System Network?
From a user's perspective, a DeviceNet network is a control system network. The ODVA DeviceNet specifications specify a DeviceNet control system network.

From a user's perspective, a ControlNet network is a control system network. The ODVA ControlNet specifications specify a ControlNet control system network.

From a user's perspective, an EtherNet/IP network is a control system network. Unfortunately, there are Ethernet networks that are extremely poor foundations for a control system network. There are also Ethernet-capable devices that don't belong on control system networks. Additionally, there are TCP/IP protocol implementations that do not recognize or acknowledge the importance of control signal data traffic and, as a result, would not handle CIP implicit message traffic appropriately.

It is clearly evident from the evidence presented in this paper that not every EtherNet/IP network is a control system network:

- An EtherNet/IP network can be constructed (but not work well) on a shared media Ethernet network.
- An EtherNet/IP network can be constructed that doesn't communicate control signal traffic (an uninteresting implementation that could not be distinguished from an information systems network).
- Any TCP/IP capable nodes on an EtherNet/IP network can considered EtherNet/IP network nodes.

Thus, the "user's perspective" qualifier is required for the premise of this paper to be a true statement. It can also be stated that given these abnormal implementations, users with that perspective need a better and more complete definition of an EtherNet/IP control system network.

The unique feature of the EtherNet/IP specifications is the specification of the implicit messaging connection for communicating control signal data. Even though the ability to communicate implicit messages is optional, when utilized, the implicit message traffic is the most critical traffic that an EtherNet/IP device communicates. The implicit messaging capability is optional only in the sense that not every device on an EtherNet/IP control system network must implement it.

How Should an EtherNet/IP Control System Network be Architected?
The EtherNet/IP specifications were developed only after switched, full-duplex Ethernet networking standards were published. In situations where there is the possibility of simultaneous traffic, Ethernet switches will always outperform Ethernet hubs. The inherent Ethernet collisions and collision recovery behaviors in a shared media Ethernet network pre-empt Ethernet QoS mechanisms. Control signal traffic is not preferentially handled in a shared media Ethernet network. There is no functional benefit to having any EtherNet/IP capable device link to a shared media Ethernet network.

The EtherNet/IP specifications require EtherNet/IP conformant devices to implement an Ethernet standard conformant communication interface. Requiring a wired EtherNet/IP conformant device to only connect to a wired Ethernet switch violates no published standard. Having wired EtherNet/IP conformant devices always connect to an Ethernet switch should be a requirement of an EtherNet/IP control system network.

In situations where there is the possibility of simultaneous traffic, full-duplex communication will always outperform half-duplex communication. There is no benefit to having an EtherNet/IP implicit messaging capable device operate in a half-duplex fashion, so why tolerate it? Ensure full-duplex communications on every cable in an EtherNet/IP control system network.

How Should an EtherNet/IP Control System Network Be Sized?
Anyone who states that a network can be virtually unlimited in size is not describing a control system network. A control system network is a local area network operating equipment connected to and controlled by the control system network. Broadcast Ethernet traffic on an Ethernet-based control system network is an Achilles heel of that network.

The network behavior in response to the broadcast list identity command[1] establishes a practical limit on the numeric size of EtherNet/IP networks. GM discovered this and shared its findings with the ODVA. It became a primary topic of discussion at the next EtherNet/IP Implementors' Workshop – the predecessor to today's EtherNet/IP Roundtable.

As a result of that experience, General Motors limited the numeric size of EtherNet/IP networks to a maximum of 250 nodes. Engineers were instructed to use two EtherNet/IP networks and exchange control traffic between them when the number of devices on a proposed EtherNet/IP network approached 220 devices.

There is a side benefit to the 250-node size limit. For networks having between 126 and 254 nodes, the first three octets of the four-octet IP version IV address can be used to readily identify the network portion of the IP address with the final octet identifying the network node number. IP networks of that size use a subnet mask of 255.255.255.0. Recommendation: For consistency and ease of training, 254 private IP node addresses should be allocated to every EtherNet/IP network, including extremely small ones.

The ControlNet Origin of the EtherNet/IP Specification Demonstrate Useful Control System Network Features
The original EtherNet/IP specification evolved from the ControlNet specifications and showed a strong desire to emulate the control signal data communication characteristics of a ControlNet network. ControlNet's producer/consumer paradigm was mimicked through the specification of multicast IP and multicast Ethernet behaviors. In retrospect, it is now generally recognized that the early EtherNet/IP

[1] Upon receiving a broadcast List Identity command from a computer just connected to the network, each EtherNet/IP device, to send the required reply to the computer, broadcasts an Address Resolution Protocol (ARP) request to determine the Ethernet address of the computer. In very large networks, this burst of broadcast traffic overloaded the capabilities of many minimally resourced devices, causing them to cease communicating. In response to GM's inquiry, the ODVA amended the conformance test suite to ensure that all EtherNet/IP devices could tolerate a burst of 250 "simultaneous" broadcast ARP messages.

specification emphasis on multicast communication of control signal data was excessive. It was unnecessarily complex and expensive to implement and support. Users discovered that most control data exchanges could be handled using point-to-point communications.

ControlNet is one of the earliest control system network technologies that has the luxury of communication bandwidth in excess of the networked control system control signal data communication needs. Prior to ControlNet, control system networks exclusively communicated control signal traffic. General-purpose information exchanges on those networks had to be embedded in control signal traffic exchanges with a device acting as the control network's gateway to the outside world.
With the introduction of an IP router for interconnecting ControlNet networks to the outside world, ControlNet both introduced TCP/IP network applications to a control system network and improved the means of acquiring information from the devices on a control system network.

ControlNet's data link protocol explicitly distinguishes control signal traffic communication bandwidth from information exchange traffic communication bandwidth. It has inherent scheduling mechanisms that ensure and enforce the behavior that information communication exchanges cannot interfere with the scheduled timing of control signal data communication exchanges. The ControlNet scheduler will not generate a network schedule that does not satisfy the message delivery timing expectations of all of the network's control signal traffic. The ControlNet scheduler will not generate a network schedule that doesn't include the ability to send at least one unscheduled message every network update period. Similar scheduling mechanisms are not part of the EtherNet/IP specifications. The best way to minimize information exchange interference with control signal data communication in an EtherNet/IP network is to use communication protocol Quality of Service (QoS) mechanisms to prioritize control signal traffic to **always** pre-empt information exchange traffic in every communication transmission queue throughout the network. However, Ethernet network collisions can disrupt Ethernet QoS behaviors. Ethernet QoS features only properly work on collision-free Ethernet networks and network segments. An EtherNet/IP control system network should always use Ethernet switches that support Ethernet QoS features configured to preferentially forward EtherNet/IP implicit message traffic.

The "One Big Switch" EtherNet/IP Control System Design Philosophy
From a user's perspective, an EtherNet/IP network is a control system network. The primary objective of a control system network is the on-time delivery of control signal traffic. To better achieve that primary objective, optimum network performance is a desirable characteristic for a control system network. The optimum performance that can be realized in a switched Ethernet-based control system network would be for all devices exchanging control signal traffic to be connected to the same line-speed, non-blocking Ethernet switch communicating in a full-duplex fashion. This observation has been named the "One Big Switch" EtherNet/IP network design philosophy.

Imagine if all the EtherNet/IP control system devices comprising an EtherNet/IP network could be connected to a single, non-blocking, line-speed Ethernet switch. Ignore any cabling costs for the imaginary switch. Distant devices could be connected to the switch by fiber optic cables. The switch needs a port for every EtherNet/IP capable device, plus a few extra ports. It should have an additional port to occasionally host a maintenance laptop computer, at least one spare port for expansion purposes, and, optionally, an extra port for connecting the switch to the outside world. See Figure 6.

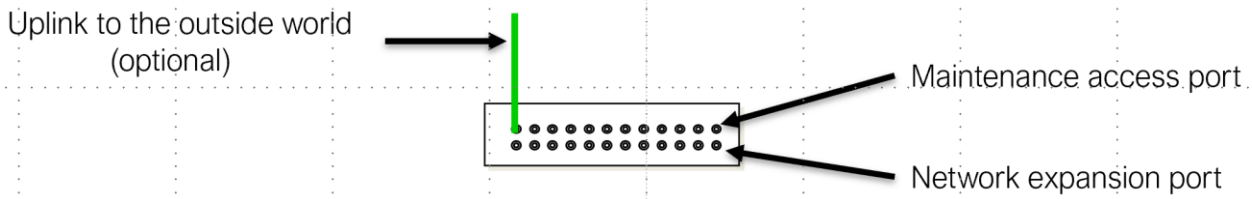

Figure 6 - Hypothetical One Big Switch EtherNet/IP Network. An imaginary, line speed, non-blocking Ethernet switch interconnecting every node in an EtherNet/IP network, establishing the optimum EtherNet/IP performance model.

All the control system traffic communicating devices in a “One Big Switch” network interact with each other across the “One Big Switch” backplane. Switch-to-switch links in a multi-switch control system network function as performance-impacting low bandwidth alternatives for what would otherwise be the switch backplane communication path of a single “One Big Switch” control system network. See Figure 7.

Control system networks control system network

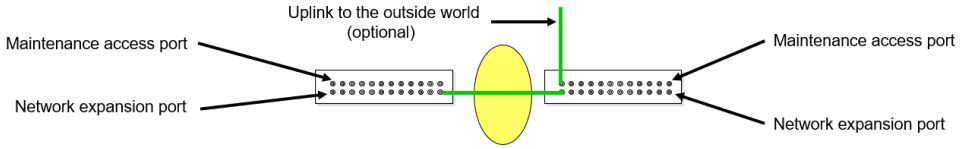

Figure 7 - Any switch-to-switch links in a multi-switch Ethernet-based control system network replace a “One Big Switch” backplane interconnection of the separated ports.

A “One Big Switch” EtherNet/IP network outperforms any multi-switch network with the same devices running the same control programs because the switch-to-switch links in a multi-switch network only emulate what would otherwise be a switch backplane interconnection in an imaginary single switch (a “One Big Switch”) network. Control signal traffic being exchanged between devices connected to separate switches will, at a minimum, experience a momentary extra switching delay that they wouldn’t experience if they were connected to the same switch. Simultaneous control traffic being exchanged between pairs of devices connected to separate switches will experience an additional traffic congestion delay. One message waits in the switch-to-switch port transmission queue while the other message is transmitting. The “One Big Switch” conceptual model establishes a switched EtherNet/IP network optimum performance target.

The only devices on a control system network that don’t communicate control signal traffic should be devices necessary to support the continued operation of the network and the control signal communicating devices on the network. Consider what happens if an Ethernet-capable device not part of the control system is connected to the “One Big Switch.” All the traffic to and from that device gets communicated through the port of the “One Big Switch” connecting the control system to the outside world. Such a device provides no benefit to and can only potentially harm the operation of the control system network. Any switch configurations required to host it are an unnecessary burden for the control system equipment maintenance personnel. The only thing it does is unnecessarily consume a port and rob bandwidth from the control system network. Such a device does not belong in an EtherNet/IP control system network. A control system network should exclusively be a control system network.

For practical reasons, only the smallest EtherNet/IP networks will be single-switch EtherNet/IP networks. Cable costs and/or a limited number of switch ports on an actual switch are both practical reasons for designing a multi-switch EtherNet/IP network. However, multi-switch EtherNet/IP networks can benefit from adhering to a “One Big Switch” design philosophy.

Ethernet switches from different suppliers and even switches from the same supplier with different firmware versions can implement nominally similar switch features in various ways. To eliminate the possibility of switch variations in a multi-switch EtherNet/IP network, all the switches in an EtherNet/IP control system network should ideally be from the same supplier’s family of Ethernet switches with a common firmware version.

To best emulate the backplane performance of a “One Big Switch” control system network, devices in a multi-switch control system network should be connected to the separate switches to minimize the control traffic being exchanged between the switches. Whenever possible, the devices that most frequently exchange control signal traffic with each other should be connected to the same switch. In other words, controllers and the devices that they control should ideally be connected to the same switch. Controllers that seldom communicate with each other should be the primary devices that are connected to different switches.

For large control system networks, the application of the “One Big Switch” Ethernet control system network design philosophy results in a switch hierarchy matching the hierarchical nature of control signal

traffic. Controllers control tools by interacting with sensors and actuators in a superior/subordinate command/response manner (Figure 8).

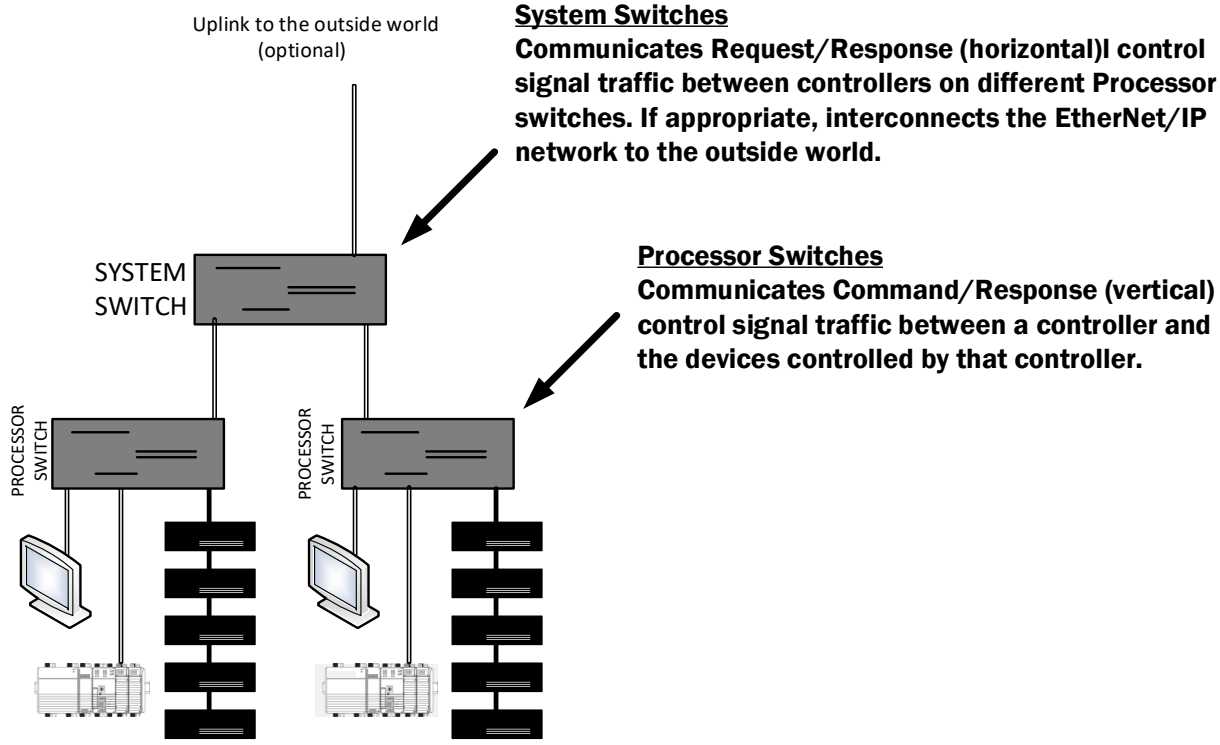

Figure 8 - Large Multi-Controller, Multi-Switch EtherNet/IP Network

Master controllers control subordinate controllers as well as sensors and actuators in a similar command/response manner. This is sometimes referred to as north/south or vertical control signal traffic. Master controllers interact with other master controllers in a peer-to-peer request/response manner. This is sometimes called east/west or horizontal control signal traffic. Maximizing the exchange of control signal traffic across the backplanes of separate switches in large multi-switch, multi-controller network results in master controllers and the equipment they control on separate switches exchanging request /response control signals between the switches. As the number of devices that a master controller either directly or indirectly controls exceeds the practical limitations of the Ethernet switch it is connected to, that switch also gets expanded in a manner that maximizes the switch backplane command/response control signal traffic and minimizes the switch-to-switch traffic with the expansion switch (Figure 9).

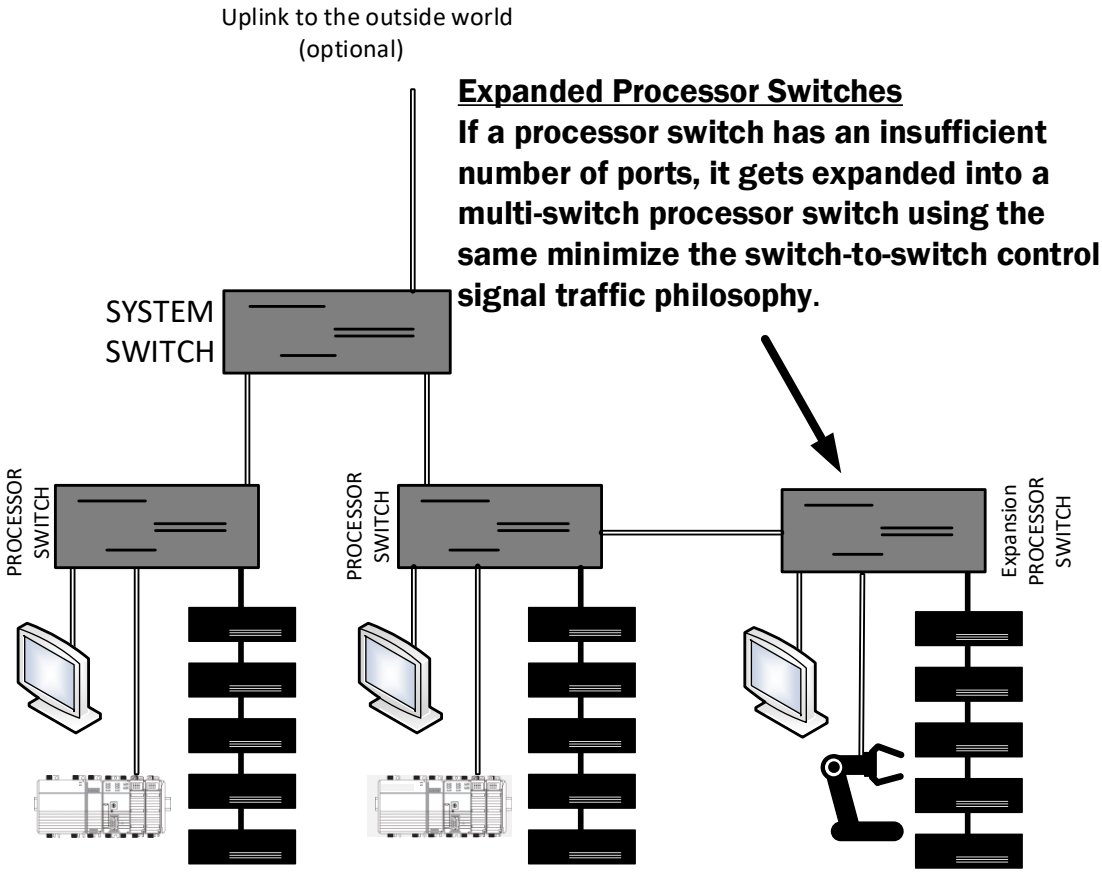

Figure 9 - Expanding a Processor Switch in a Large Multi-Controller Multi-Switch EtherNet/IP Network

Using I/O Networks vs Controller Networks
Some end users, such as General Motors, use DLR devices primarily in linear EtherNet/IP DLR I/O network segments connected to the traditional switch. This architecture saves the cost of extra switch ports and the cost of an additional Ethernet interface card for the controller.

A significant drawback to recognize when using linear DLR I/O network segments is that the loss of power to a device in a segment disrupts any communication with the downstream devices on that network segment.

One particularly useful application is to use a DLR I/O device as a copper cable extender. It functions as an easily monitored, extremely cost-effective unmanaged switch that allows devices to be copper cable connected to traditional switches up to 200 cable meters away.

Recommendation: Evaluate using DLR in linear I/O segments.

Connecting an EtherNet/IP network to the outside world
A crucial aspect of a “One Big Switch” Ethernet based control system network is how a user chooses to utilize the optional uplink to the outside world port. More data is now required from manufacturing machinery and much of that data originates in devices participating in the production system. To achieve the visions of Industry 4.0 and Smart Manufacturing, the ways in which control networks and information networks interact are becoming more frequent and more complex than ever. That presents some difficult challenges and numerous new options for users architecting production systems using control system networks like EtherNet/IP.

There are numerous control system networks that don't connect to the outside world. The cybersecurity and network support capabilities of isolated EtherNet/IP networks would seem to be the base foundation upon which to build and expand when they get connected to an Information Technology (IT) network. From a user perspective, the network support for these types of isolated EtherNet/IP control system networks are not being adequately addressed by the EtherNet/IP specifications. Perhaps cybersecurity specifications for ControlNet networks should be developed, allowing them to serve as the model for a cybersecurity foundation for isolated EtherNet/IP networks. Network support tools appropriate for both the industrial environment and the control system tooling support are persistent user needs.

There are other EtherNet/IP networks that only connect to the outside world by using an application gateway. These EtherNet/IP networks also don't have access to information system network management or cybersecurity servers. Again, from a user perspective, the network support and cybersecurity needs of these types of EtherNet/IP control system networks are not adequately addressed by the EtherNet/IP specifications.

In the isolated and gateway interconnected EtherNet/IP networks, control signal traffic is confined to the control system network. Datalink bridges can be used to pass control signal traffic between two devices on separate, isolated control system networks, but they always experience a noticeable delay when relaying control signal traffic. The most common form of datalink bridge is essentially two input /output modules with a shared memory acting as a recognized device on both networks. Control signal traffic sent as an output signal from a device on one network can be read as an input signal by a device on the other network and vice versa.

The Tremendous Impact of a Single Cable

The ControlNet networks that GM initially installed were gateway interconnected to the IT plantwide Ethernet information network. See Figure 10.

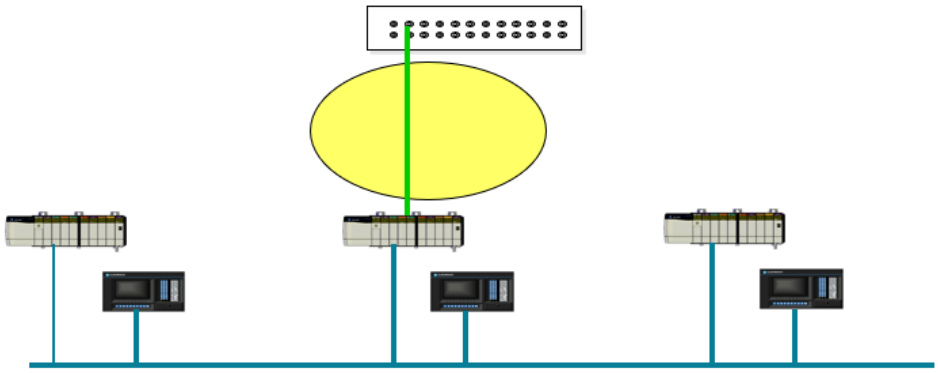

Figure 10 - Ethernet cable linking a ControlNet network to the outside world using an Ethernet to ControlNet gateway

A cable ran from an IT access layer Ethernet switch to one of the ControlNet capable PLCs programmed to be the ControlNet network gateway. The initial EtherNet/IP networks installed at GM were direct replacements for control system networks that would have alternatively been ControlNet networks.

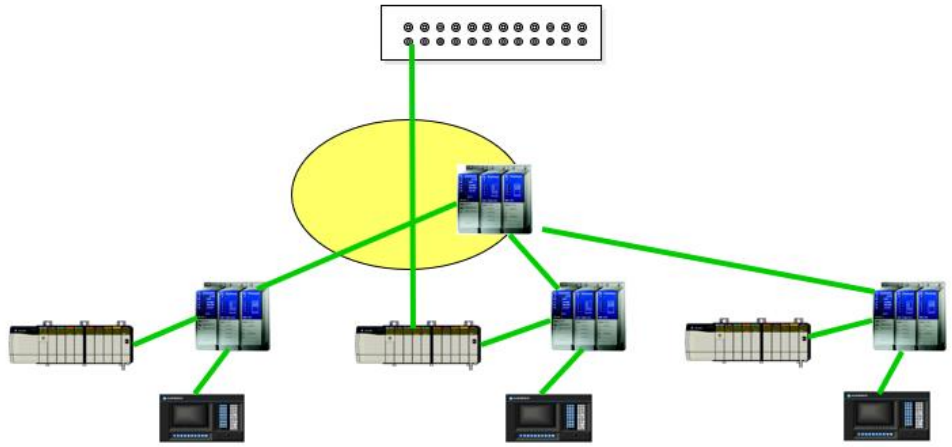

Figure 11 - Ethernet cable linking an EtherNet/IP network to the outside world using an Ethernet to EtherNet/IP gateway

One PLC could have been programmed to act as an Ethernet to EtherNet/IP network gateway (See Figure 11). However, the decision was made to route the cable from the IT network switch that normally ran to the control network gateway PLC to one of the EtherNet/IP network switches instead (See Figure 12).

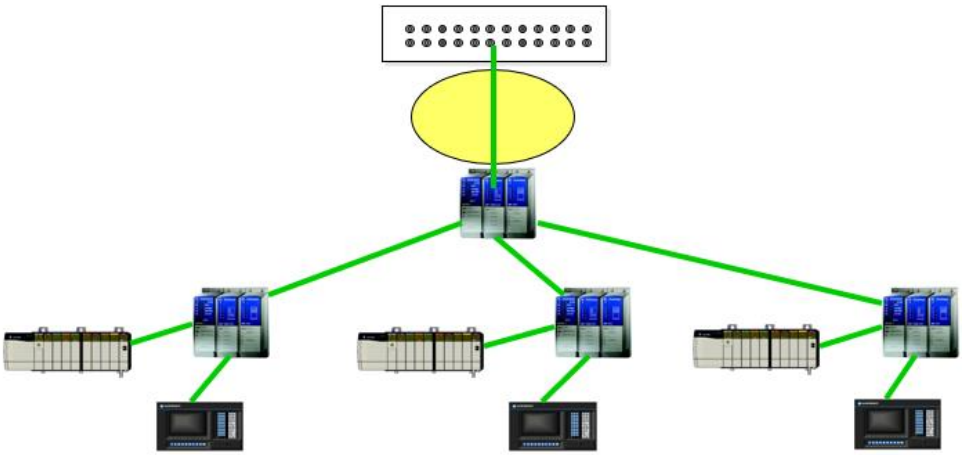

Figure 12 - Ethernet cable linking an EtherNet/IP network to the outside world using a direct switch-to-switch link.

The re-routing of that single cable exacerbated multiple IT/OT culture gap issues. What used to be two entirely segregated networking environments with unaligned objectives and philosophies were abruptly transformed into a logically integrated internetwork. Independent organizations that normally never interacted and ignored each other were immediately required to cooperate and coordinate. What used to be separate policies, procedures and practices for network design, network addressing, hardware procurement, network installation, network security, network support, and network change control now needed to be coordinated and harmonized.

Directly Interconnecting Ethernet Infrastructure Devices Transforms Two Previously Separate and Distinct Networks

With the development of line-speed Ethernet network to Ethernet network routers (AKA Layer 3 switches) it became feasible to route EtherNet/IP implicit message traffic between separate Ethernet networks and between VLANs within an Ethernet network. An EtherNet/IP Specification Enhancement was passed allowing unicast EtherNet/IP implicit message traffic to be routed between Ethernet networks and Ethernet VLANs.

The ideal way to interconnect an EtherNet/IP network with an IT information system network is by configuring a “one big switch” EtherNet/IP network to be a VLAN of the IT TCP/IP Ethernet information

system network. Designating the EtherNet/IP network as an IT information system network VLAN segregates all broadcast and multicast Ethernet traffic of the separate networks and forces all traffic to and from the EtherNet/IP network to be routed by an information system network router. Additional IT network security features can be implemented at the router. If the IT network is being used to exchange EtherNet/IP implicit message traffic between EtherNet/IP networks, any IT network routers handling EtherNet/IP implicit message traffic should be line-speed layer 3 switches.

Layer 2+ Ethernet switches were enhanced with a static routing feature, allowing for limited inter-VLAN traffic static routing. GM uses that feature to exchange traffic between EtherNet/IP controller-level and EtherNet/IP I/O level networks. VLANs are used to segregate any broadcast or multicast traffic on different EtherNet/IP networks connected to the same static routing capable switch. EtherNet/IP I/O network traffic is constrained to the switch. The only devices that aren't members of the EtherNet/IP I/O level network that can communicate with the devices on the I/O level network are the devices on the EtherNet/IP controller-level network that have awareness of the static routing capability of the Layer 2+ routing capable switch.

Unicast control signal traffic can be routed between multiple devices on different EtherNet/IP networks connected to an information system network with virtually no delay simply by routing that traffic through a layer 3 switch on that information system network.

Summary

This paper started by posing two fundamental questions: "What is an EtherNet/IP network?" and "What is a control system network?" The current definition of an EtherNet/IP network was found to be both nebulous and more expansive than the definition of a control system network should be.

While a control system network is a network of computerized devices with the primary purpose of exchanging control signal traffic, there is no such purpose included in the current definition of an EtherNet/IP network. Analyzing ODVA documents and exploring various kinds of EtherNet/IP network configurations leads to an inference that any Ethernet based IP subnet with at least two EtherNet/IP conformant nodes forms an EtherNet/IP network. And more unpalatable, there is no way to refute the claim that any other node on that subnet is an EtherNet/IP network node. Simply by virtue of the IP address assigned to them, devices that are not EtherNet/IP capable would become EtherNet/IP network nodes. A managed switch that only hosted EtherNet/IP conformant devices from the same EtherNet/IP network would not be a member of that EtherNet/IP network if was configured with an IP address from a different subnetwork range.

This paper claims that, from the user's perspective, an EtherNet/IP network is a control system network. To remove the "user's perspective" qualifier from that claim requires changing the question "Is an EtherNet/IP network a control system network?" into the unqualified statement "An EtherNet/IP network is a control system network". Doing that would require a definition clarifying that EtherNet/IP networks exist to exchange control signals. TCP/IP Ethernet networks that don't exchange control signals are not control system networks, even when they have EtherNet/IP capable devices on them.

Only devices engaged in exchanging control signals and those devices necessary to support the devices that exchange control signals belong on an EtherNet/IP control system network. End device membership in an EtherNet/IP control system network should be the same whether or not that network is connected to the outside world. The only device to be added to an outside world interconnected EtherNet/IP control system network should be the router used by devices in the network to communicate with the outside world.

In that context, this paper introduces the "One Big Switch" optimum performance model for an Ethernet based control system network and provides several recommendations for architecting performance oriented EtherNet/IP control system networks.

Finally, the paper discussed the interconnection between EtherNet/IP networks and plant Information networks. That interconnection can transform previously separate and distinct networks into a single internetwork with wide-ranging and unanticipated consequences.

Gary Workman
Professional consultant located in Metro Detroit focused on user-oriented perspective of networking solutions. Retired from General Motors as Principal Engineer for Plant Floor Networks. More than 48 years of experience, with 35 in the networking field bridging the divide between Information Technology and Operational Technology.

During Covid years co-wrote the book “The Everyman’s Guide to Properly Architecting EtherNet/IP Networks” available from Amazon.com with John Rinaldi.

MBA from Harvard Business School

John S Rinaldi
John Rinaldi is Chief Strategist and Director of WOW! for Real Time Automation (RTA) in Pewaukee WI. With a focus on simplicity, support, expert consulting and tailoring for specific customer applications, RTA is meeting customer needs in applications worldwide. John is not only a recognized expert in industrial networks and an automation strategist but a speaker, blogger, the author of more than 500 articles on industrial networking, and six books including Industrial Ethernet, OPC UA: The Basics, Modbus, OPC UA - The Everyman's Guide and ETHERNET/IP.

John Rinaldi
Real Time Automation
http://www.rtaautomation.com/contact-us/
https://www.linkedin.com/in/johnsrinaldi