Troubleshooting EtherNet/IP Networks

Merrill Harriman
Network Systems Architect
Schneider Electric

Presented at the ODVA
2014 Industry Conference & 16th Annual Meeting
March 11-13, 2014
Phoenix, Arizona, USA

Abstract:

As the complexity of our customer's industrial Ethernet networks continues to increase, so does the challenge to diagnose problems in these complex networks. Yet the skill set of the typical automation maintenance person is not increasing at a comparable rate. This leaves a void in the resources required to keep these systems running. Our customers need help with understanding how to utilize the diagnostics capabilities of the products and how to interpret the underlying meaning of this data.

This paper will discuss some of the more common physical layer and data link layer problems in EtherNet/IP™ networks and trace from symptoms to most likely root causes. This paper discusses the underlying meaning of some of the more esoteric diagnostic counters and how they can be used for practical troubleshooting. Profiles of network problems will be presented to help zero in on a root cause. This paper focuses on two key diagnostic data accesses: EtherNet/IP diagnostics objects, and SNMP (Simple Network Management Protocol). A brief description of the underlying technologies will be given but it is assumed that the reader has some familiarity with concepts involved. While not intended to be all inclusive of every potential network problem, the concepts and techniques presented here provide a basis for how diagnostic data can be interpreted and used.

Keywords:

Troubleshooting, Diagnostics, Performance, Link Object, SNMP, Counters, EtherNet/IP, Ethernet, Network, Duplex Mismatch, Electrical Noise

1. Introduction

Often when facing a network problem, network engineers and maintenance personnel will gravitate towards Ping and packet capture as their two favorite tools. But these are not always the most effective or efficient in diagnosing low level network problems. Diagnostic data is available to help diagnose these problems but the data is often not well understood nor do people understand how it can be used. This paper will address each of these concerns.

The paper begins with a brief discussion of the pros and cons of a few key sources of diagnostic data. It then introduces the set of detailed diagnostics that is available from end devices via EtherNet/IP and from infrastructure devices via SNMP (Simple Network Management Protocol). It also shows how they are relevant to an EtherNet/IP network. The paper also shows a correlation of the data from each access method. It then applies this data through case studies to show how they can be used to diagnose some common network problems. Finally, it provides a complete cross reference of diagnostic data from EtherNet/IP to SNMP.

2. Glossary of Terms and Acronyms

Term / Acronym	Definition
Autonegotiation	A communications algorithm that provides negotiation of the duplex setting used by devices on each end on an Ethernet link
EtherLike-MIB	Defines objects that are specifically relevant to Ethernet based devices
Full Duplex	A communications mode that allows bidirectional simultaneous communication
Half Duplex	A communications mode that requires unidirectional communication
MIB	Management Information Base – A file defining a collection of management data objects
MIB Browser	A software package that utilizes SNMP to access MIB data objects from a device
MIB II	A collection of fundamental data objects as defined in RFC 1213
OID	Object Identification – An addressing schema used to identify MIB data objects
PC	Personal Computer
SNMP	Simple Network Management Protocol

TABLE 1 - TERMS AND ACRONYMS

3. Overview of technologies used in troubleshooting

3.1. Web

Many industrial control devices support some form of web pages whether it be for product information, configuration, diagnostics, or some other form of page. Each vendor has their own approach to their web content. Most devices, however, will support some form of network diagnostics web pages and the content of these pages often reflects the content of the EtherNet/IP Ethernet Link Object (0xF6). This is not governed; it is more of a de facto standard.

For this paper, most emphasis will be placed on the EtherNet/IP objects and SNMP. If the devices in the network do provide web pages with this same information it may be easier to access. The theories involved still apply. But since web support is not guaranteed to be consistent, for this paper it is not assumed to exist.

3.2. Ping

Ping, also known as ICMP Echo Request/Reply, is one of the most fundamental and ubiquitous diagnostic tools. It sends ICMP echo requests to the target device and maintains statistics on connectivity and response times. A Ping request can be initiated from the Windows PC command line as follows:

```
C:\>ping 192.168.123.1
Pinging 192.168.123.1 with 32 bytes of data:
Reply from 192.168.123.1: bytes=32 time=1ms TTL=64
Reply from 192.168.123.1: bytes=32 time=1ms TTL=64
Reply from 192.168.123.1: bytes=32 time=1ms TTL=64
Reply from 192.168.123.1: bytes=32 time=1ms TTL=64

Ping statistics for 192.168.123.1:
   Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
   Minimum = 1ms, Maximum = 1ms, Average = 1ms
```

FIGURE 1 - PING

Ping is one of the most often used diagnostic tools. First and foremost it is used for determining "reachability" to see if two devices in the network can talk to each other. This can be very effective at helping to decipher the cause of a connection establishment problem between an application problem or a network configuration or fault. If the device can be pinged then likely the device is also able to receive connection requests and the network is not the likely cause.

Ping is also one of the most over used diagnostics tools. It has its place as a diagnostics tool but it is not very effective at diagnosing most performance related problems in industrial control systems. It will give an indication of response times and this is very effective when diagnosing complex networks with multiple routing layers. But for relatively simplistic networks where response times are typically very fast, the PC's discretion in this time measurement is not adequate for determining network performance delays. Much of the lack of accuracy is attributable to lack of resolution in the time stamping in the M.S. Windows stacks. Ping is also not a good indicator of device performance delays caused by a very busy device because the Ping processing happens at a much lower level in the device firmware than does an application protocol such as EtherNet/IP.

Ping can give some indication of problems in a very high packet loss environment. But these types of problems are often more statistical in nature and affect ongoing packet flows. Unless the situation is very dire the probability that the particular ICMP packet will be affected may be slight. Therefore, a quick ping check may indicate that the network is fine but in fact there are problems. The ICMP packets may have snuck through without getting clobbered. On the contrary, if only some of the ping requests fail then it is a good indication that the network is in trouble.

If all of the ping requests fail then it could be blocked by a firewall, disabled at the target for security, a network configuration/wiring problem, a device configuration/fault problem, or an extremely high noise or high bandwidth utilization problem. A failed ping request does not provide many answers.

3.3. Packet Capture Tools

A packet capture tool is another very ubiquitous tool in network troubleshooting. It will copy packets from a PC's interface into a file and parse these packets for analysis by the user. This file is often known as a "sniffer trace". These tools packet parsing and filtering capabilities make them indispensible when troubleshooting very complex network problems. They allow detailed timing analysis and will dissect each packet in intimate detail.

It is because of this extreme detail that packet capture is another often overused tool. It is all too easy for the troubleshooter to get lost in all this detail. They can easily overlook the big picture or be fooled into chasing the wrong problem or a non-existent problem and wasting a significant amount of time. They must also understand that the location in the network from which they take the packet capture can be critical to getting useful information. In some situations, such as corrupted packets, a packet capture may not provide the complete information to allow diagnosis of the problem. It is important to understand when packet capture is applicable.

3.4. EtherNet/IP Diagnostics Tools

The EtherNet/IP protocol defines diagnostics objects that can provide invaluable information for troubleshooting network problems. Access to these objects can be obtained through a PLC, SCADA, or various software packages that run on a PC. Some vendors offer free tools that can query a device for these diagnostics objects. These objects are one of the primary focuses of this paper and will be discussed in detail throughout. This paper does not recommend any particular tool or method of access and assumes that the reader has a solution available.

3.5. SNMP MIB Browsers

As will be discussed later, SNMP provides a wealth of information and is the primary protocol used with infrastructure devices for diagnostics and configuration. SNMP is also supported on some EtherNet/IP industrial automation devices. Software packages known as “MIB browsers” (Management Information Base) utilize the SNMP protocol to provide access to this information. Numerous vendors provide MIB browsers and many are free tools readily available for download from the internet. As well, SNMP can be used with some SCADA systems and with network monitoring tools.

Some MIB browsers provide extended capabilities such as automatic polling of data and graphing of multiple data objects over time. Another common feature is properties dialog that provides detailed information of each object including its data type and a textual description. Many will also translate obscure numerical state values into human readable strings to aid with the understanding of the value of the object. SNMP MIB browsers are in indispensible tool especially when working with infrastructure devices.

3.6. Network Monitoring Tools

Network monitoring tools will periodically poll diagnostic data to help detect network problems before they become severe and provide detailed diagnostic information to help guide the troubleshooting efforts. These tools typically rely heavily on SNMP and will analyze the diagnostic data against alarm conditions to provide an event log and to send alerts when problems are detected. Some will also monitor for changing conditions in the network such as configuration changes that may be the root cause of the pursuant problems. Network monitoring tools can provide vital information to help optimize troubleshooting efforts and reduce network down time.

4. Network Diagnostic Data Access

This paper will focus on two methods of accessing diagnostic data; EtherNet/IP and SNMP. Web access will not be a primary focus as different devices may have significantly different levels of support for diagnostic information via their web pages. EtherNet/IP and SNMP support is, empirically, more consistent from device to device. That is not to say that web pages are not an important source for easy access to diagnostic information, it is simply that generic discussions favor generic solutions.

It is assumed throughout this paper that all end devices will support EtherNet/IP diagnostics objects and that all infrastructure devices will support SNMP. This further assumes that all infrastructure devices are, quite appropriately, managed devices. In some cases the infrastructure devices may support both SNMP and EtherNet/IP but that is not assumed in this paper.

There is some ambiguity that arises with multiport end devices. Should they be considered an infrastructure device? It is assumed that these support EtherNet/IP diagnostics as they are end devices. Beyond that, concerning SNMP support, this paper does not assume one way or the other.

4.1. EtherNet/IP Network Diagnostics Objects

The EtherNet/IP Ethernet Link Object (0xF6) may appear daunting with numerous esoteric counters but it provides a significant amount of useful diagnostics information. Part of the objective of this paper is to demystify the meaning and usage of this diagnostic information. Later sections will explain each of these attributes and show how they apply to EtherNet/IP networks. These attributes will then be discussed in the realm of various network problems and show how the attributes can be used to shed light on the situation.

The EtherNet/IP TCP/IP Interface Object (0xF5) is more focused on configuration properties than diagnostic counters; but the information contained here can be vital to understanding a problem on a network. Misconfiguration is a leading cause of network problems and many of the clues to those problems can be found here.

4.2. SNMP

Simple Network Management Protocol (SNMP) provides access to a wealth of information from most infrastructure devices and many end devices. It uses files that contain a predefined set of data objects called Management Information Base (MIB) files. These files group together common objects and each object definition contains a complete description of the object including the data type, semantics, description and other such properties.

Through a relative addressing schema, each object in the MIB file is assigned a unique hierarchical Object Identification (OID). The OID forms a dotted decimal string that could be likened to a path to the object class and further to define an instance of the object. Each unique object also has a unique name that is derived from its location in the MIB tree hierarchy and the data object it represents. For example, "ifDescr" is a unique object that provides a textual description of an interface on a device. It is found in the interface table at the OID = (1.3.6.1.2.1.2.2.1.2) which is defined in RFC 1213 MIB II. The full hierarchical path to the object is shown here with the textual name and associated OID: "iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).interfaces(2).ifTable(2).ifEntry(1).ifDescr(2)"

There are numerous standard MIBs available, each defined in a corresponding RFC. MIB II is the most fundamental of these with a number of "mandatory" objects. Essentially any device that supports SNMP will support some if not all of MIB II. Beyond this, many support additional other standard MIB files such as the "EtherLike-MIB" that defines objects that are specifically relevant to Ethernet based devices. Additionally, many advanced devices, especially infrastructure devices, support proprietary information via a "private MIB" a.k.a. "enterprise MIB". These MIBs allow the vendor to define their own set of data objects that can be accessed via SNMP.

Software tools known as MIB browsers will load these MIB files and typically present the objects in a natural hierarchical folder relationship. The MIB browser will allow the user to query a device for each object and present the response. The query message to the device contains the OID of the object and instance of the value being requested. The device parses this OID to process the request and returns the response. If the device does not support the object then it will return an error indicating not supported. The MIB Browser will then parse the response against the definition in the MIB file and present the result to the user.

5. Interpreting Diagnostic Counters

5.1. Error & Statistics Counters

There are a number of error and statistic counters that can be essential to network troubleshooting. Some of these are found in the EtherNet/IP Ethernet Link Object (0xF6). These can also be found through SNMP. For a complete correlation of EtherNet/IP and SNMP error and statistics counters, (see **EtherNet/IP Diagnostics and SNMP Correlation**) section of the index.

With the Ethernet Link Object (0xF6) these counters are provided in two separate complex attributes; Interface Counters (Attribute # = 4) and Media Counters (Attribute # = 5). The Interface Counters maintain a direct correlation to counters found in "RFC 1213 MIB-II Management Information Base".

The Media Counters, as defined in "The CIP™ Networks Library Volume 2 EtherNet/IP Adaption of CIP" are derived from "RFC 1643 - Definitions of Managed Objects for the Ethernet-like Interface Types". This definition of the counters has been reclassified to historical status by "RFC 3638 - Applicability Statement for Reclassification of RFC 1643 to Historic Status". In its place since September of 2003 is a new definition of the counters in "RFC 3635 - Definitions of Managed Objects for the Ethernet-like Interface Types". This is the implementation found in most current infrastructure equipment. For the counters used throughout this document, the definitions have not changed significantly. The SNMP based counters and the associated EtherNet/IP counters are assumed to be equivalent.

For the purposes of this discussion, these counters are separated into two different groups. One group deals with counters associated with packet collisions. These will be discussed in the next section. The remainder of these counters is discussed here to help clarify the meaning behind the counters.

Counter	Definition
In / Out Ucast Packets	The number of unicast packets received / sent on the interface. For EtherNet/IP this will include all point-to-point connections and explicit messaging connections.
In / Out NUcast Packets	The number of non-unicast packets received / sent on the interface. This includes all broadcast (ARP, DHCP, Bootp, NetBIOS, …) and multicast traffic. For EtherNet/IP this includes multicast class 0 & 1 connections.
In / Out Discards	The number of good packets that were discarded; typically because the input / output queue was full. (See **Bad Performance Caused by Overloaded Devices**).
In / Out Errors	This is an aggregation of a multiple specific errors. (See Table 3 - Duplex Mismatch Counter Correlation).
FCS Errors	Frame Check Sequence errors – a form of CRC that validates the packet against corruption.
SQE Test Error	This test is designed to fix a problem in earlier versions of Ethernet. With today's hardware it is seldom if ever used. If this counter is counting the device may be failing.
MAC Transmit / Receive Errors	Frames for which transmission or reception fails due to an internal sublayer transmit/receive error. The actual definition of this counter is implementation specific but typically this counter covers any error that is not covered by any of the more specific counters. If this counter is counting the equipment may be failing.
Carrier Sense Errors	On a half duplex link carrier sense errors are to be expected. If this counter is counting on a full duplex link the device may be failing[1].

TABLE 2 - NETWORK DIAGNOSTIC COUNTERS

Reference 1: https://supportforums.cisco.com/docs/DOC-1806

5.2. Collisions and Related Counters

At first blush, the reader may be tempted to overlook this section as their network, like most modern networks, is using full duplex; thus, they assume that inherently collisions will not be a problem. But, due to the way the duplex negotiation process works, half duplex can happen "by accident" and it is important to understand the ramifications.

Ethernet networks that use twisted pair copper with 10 or 100 Mbit/s connections (10Base-T and 100Base-TX) utilize 2 twisted pairs of the available 4 twisted pairs for transmission and receive. Modern Ethernet cables are configured for cross-over and one pair carries packets in one direction and the other pair carries packets in the other direction. In legacy networks that utilize hubs and repeaters these pairs formed a virtual bus and the devices share the network. Therefore, only one packet could be on the wire at any given time. Collisions occur when two (or more) devices in a shared network attempt to transmit at the same time. Ethernet hubs extend the shared network to all connected devices and effectively extend the collision domain to all devices in the shared network.

Ethernet switches used with half duplex links restrict the collision domain to only the two devices on the link; they can still have collisions but it reduces the probability. In actuality, with switches no real collision occurs. Since each direction has its own pair the network signals never actually "collide". But to be backward compatible, a device must limit the communication to one direction at a time. To do this it listens on it's receive pair while it transmits on it's transmit pair. If during the transmission it detects a transmission from the remote end on it's receive pair it will declare a collision has occurred and activate its collision circuitry.

In a half duplex link, collisions are a normal occurrence. But these should be detected early enough in the transmission by the hardware to allow a back-off and retry mechanism to renegotiate the transmission. The key is that the collision must be detected within the first 64 bytes of the transmission. This is the minimum length of an Ethernet frame.

When an Ethernet device utilizing a half duplex link attempts to transmit a frame it will use its Carrier Sense Multiple Access / Collision Detection (CSMA/CD) circuitry. It will sense the line for the presence of other traffic and if present, it will defer its transmission for a short time and try again. When it does this it also increments a hardware counter called "*Deferred Transmissions*". This counter provides an indication of link utilization as a very busy link will cause this counter to increment. This counter is only applicable to half duplex links. By definition the full duplex link will never be busy when the device attempts to transmit as it does not wait for a clear line.

If the half duplex connected device sees a clear line and begins its transmission at the same instant another device also begins, there will be a collision. The time it takes to detect the other devices transmission is associated with the propagation delay of the other devices signal through the wire. When the collision is detected, the device will cease its transmission of the current packet and instead will immediately begin sending a "jam" signal on the wire to inform the remote end of the collision. It also increments its internal collision counter. It will then back off for a random delay and retry the transmission. It will continue to retry until one of the following cases occur:

1. If after the single collision it successfully transmits the packet it increments the "*Single Collisions*" counter.
2. If after less than 16 collisions it successfully transmits the packet it increments the "*Multiple Collisions*" counter.
3. If after 16 collisions it still does not successfully transmit the packet it will discard the packet and increment the "*Excessive Collisions*" counter.
4. If it starts the transmission and detects a collision after the first 64 bytes of data are already transmitted it will stop the transmission, discard the packet, and increment the "*Late Collisions*" counter.

All of these counters are only applicable to half duplex links. In a full duplex system they should never increment. If these counters are counting in a presumably full duplex link, look for a possible duplex mismatch. This topic is covered in detail later in this paper.

5.3. Raw or Rate

Diagnostic counters, in their raw form as they would be retrieved via SNMP or the EtherNet/IP Diagnostics Objects, present monotonically increasing count of items since the module was last booted. The counter by itself does not provide any evidence as to when in the past the count increased; it may have been in the past year, the past week, or the past minute. Raw counter values can be misleading. The important point is not to be lead astray by a high numerical value; it may have no relevance to the current problem.

To get a true sense of the diagnostic data's relevance to the current situation, it is usually best to look at the rate of change rather than the actual value. This can be done manually by simply noting the value and polling the value again in a few seconds. Or some diagnostic tools including MIB browsers will plot the value against time. Many network monitoring tools provide logging reporting capabilities to allow historical analysis of the data. Additionally, data can be sampled to a file and plotted with M.S. Excel or some other plotting tool.

6. Network Problems that Cause Bad Performance

Bad performance in an industrial control system is a subjective term but it is typically characterized by longer than expected response times in the system. These excessive response times could manifest as slow reactions in the physical world (a gate is slow to react), or could manifest as slow updates on a SCADA screen for example. Whatever the case, the user has reason to believe that the system is slower than they expect.

Of course the bad performance could be an application problem rather than a network problem. The two classes of problems are often cross-diagnosed and what is thought to be an application problem is in fact a network problem or vice versa. For the remainder of this discussion it is assumed that the bad performance is correctly diagnosed as a network problem.

The following sections will outline some typical network faults and detail how they present in an industrial automation network. In each case the background of the problem is described with a description of how it affects performance and some methods of diagnosing the problem are given. In some cases the clues to discerning one root cause from another may be subtle. The reader is advised to understand each case before jumping to a conclusion.

6.1. Understanding Packet Loss

Packet loss, for this discussion, is a random loss of a packet from an ongoing stream of packets. Packet loss can be the result of a number of causes, some of which are discussed in detail in this document. Packet loss can be the result of packets being dropped by devices that are overloaded, packets being corrupted by electrical noise or faulty equipment, configuration errors, or internal errors in a device. Whatever the root cause, packet loss will have similar symptoms at the transport and application layers.

For TCP based connections, EtherNet/IP class 2 or 3 for example, as packets are lost, the TCP connection will fall back to TCP retries to maintain the connection. This can be seen via SNMP in the tcpRetransSegs (1.3.6.1.2.1.6.12) {RFC 1213}. This information is not available via the EtherNet/IP diagnostic objects.

EtherNet/IP class 0 or 1 connections use UDP and the notion of the connection is maintained at the application layer. Since UPD is a connectionless protocol, the packets are dropped with no obvious indication in a counter. For UDP based connections, EtherNet/IP class 1 connections for example, there is no inherent counter that indicates that messages have been missed in the stream – such as an indication of non-sequential sequence numbers. Some devices do support this via their web pages but it is not part of the current diagnostics objects.

From the transport layer or the network layer, the root cause of the problem is not always evident. Problems involving some form of packet loss require a more detailed look at what is happening in the lower layers of the network communications. This is the role of the Ethernet Link Object and the associated SNMP objects. The following sections will discuss a selection of network problems and show how these can be diagnosed with a detailed look at the lower layer diagnostic information.

6.2. Bad Performance Caused by Duplex Mismatch

Often industrial automation network engineers will opt for statically configured speed and duplex settings for the ports of their industrial Ethernet equipment. And often they opt for a full duplex connection. However, most vendors will provide a default port configuration of “Autonegotiate” (AutoNeg). Autonegotiation attempts to automatically resolve an agreed upon duplex of a link by communicating with the link paired equipment (the device on the other end of the link) to negotiate a mutual set of values. The speed setting is detected by a link pulse signal sent from the link pair device and is not part of this negotiation.

When a statically configured device is connected to a device configured for AutoNeg, the statically configured device will not participate in the link negotiation process. However, it will still send “link pulses” and from these the AutoNeg device can determine link speed but it cannot determine the duplex configuration of the statically configured end. As per IEEE 802.3, if a device cannot determine that the remote end can support full duplex, it must assume half duplex. The AutoNeg device then configures itself for the correct speed, maybe 100 MB, and assumes half duplex. Yet the statically configured end may have been configured for 100 MB / full duplex. Since the speed setting is correct, the two devices will be able to understand each other’s transmissions and the link LEDs will show as expected. But, the mismatched duplex settings will likely cause errors in the communication stream. Simple slow data transfers (ping, for example) may work without error but as traffic on the link increases the probability of error also increases.

Problems arise when both devices attempt to send data at the same time. If the AutoNeg half duplex end attempts to transmit while the other end is transmitting it will detect a collision of the two transmissions. The autoneg half duplex end uses CSMA/CD (Carrier Sense Multiple Access / Collision Detection) that will then detect the collision, back-off for a random delay and retry the transmission. This all takes place at the hardware level in the MAC layer – at the layer 2. This happens regardless of the transport protocol of the packet; it is not a TCP retry.

On the statically configured full duplex end, the device assumes that collisions are not possible (it can receive and send at the same time). The CSMA/CD collision detection circuitry is disabled. Hence, there is no associated collision detection, back-off, and retry algorithm on this end of the link. Since it does not even detect the collision and retry the packet, any packets from this end resulting in collisions are lost. With this high rate of packet loss, the performance of a duplex mismatch link can be substantially worse than a properly configured half duplex link.

As traffic increases the probability of two coincident transmissions also increases. Since it is a problem of probabilities, detecting this situation with a simple ping test that only transfers a few packets may be difficult. The link appears to be fine except it is randomly dropping packets. This problem is also often incorrectly attributed to noise on the Ethernet cable, a bad cable or a failing switch or end device.

Typically these situations are finally detected by manually checking the configuration of the affected equipment. But one must know to look for this particular configuration error. Actively monitoring for a duplex mismatch situation with PLC logic or a network management tool can also be used. This is one place where EthereNet/IP diagnostics has an upper hand over SNMP diagnostics as the duplex information is not available via SNMP on most end devices. The EtherNet/IP Ethernet Link Object (0xF6) through the Interface Flags attribute does provide both the configured autonegotiation setting and the operational duplex status. Additionally, it has a 2 bit status code in the Interface Flags to warn of failed autonegotiation attempts. With SNMP, support for the duplex setting and duplex operational state objects can be somewhat elusive and is often only available on proper managed switches.

Monitoring for duplex mismatch explicitly does require a component in the system to be configured to actively check this information routinely. This becomes more challenging when an EtherNet/IP device that does not support SNMP is connected to a switch that does not support EtherNet/IP.

If you know what to look for, however, you can detect this situation empirically. Since on the statically configured full duplex end of the link the collision detection (CSMA/CD) has been disabled, the associated collision counters will also not get incremented. Yet on the opposite end of the link, the end with the autonegotiation configured and is operating in half duplex, the counters will be incrementing. For a properly configured half duplex link to detect collisions at only one end of the link is a statistical improbability as it implies that the jam signal is also never detected. Therefore, any link that is detecting collisions on only one end is very likely misconfigured for duplex.

Additionally, the full duplex end may be detecting FCS Errors (Frame Check Sequence) and/or Alignment Errors on its received packets. This is due to the half duplex ends transmitted packets being truncated to send the jam signal indicating collision detection. As these corrupted packets are received at the full duplex end of the link they are detected to have FCS errors or alignment errors. On the half duplex end of the link the FCS and Alignment Errors will be significantly less since the full duplex end never asserts the jam signal during its transmission and the packets are received without error or discarded as the result of a collision.

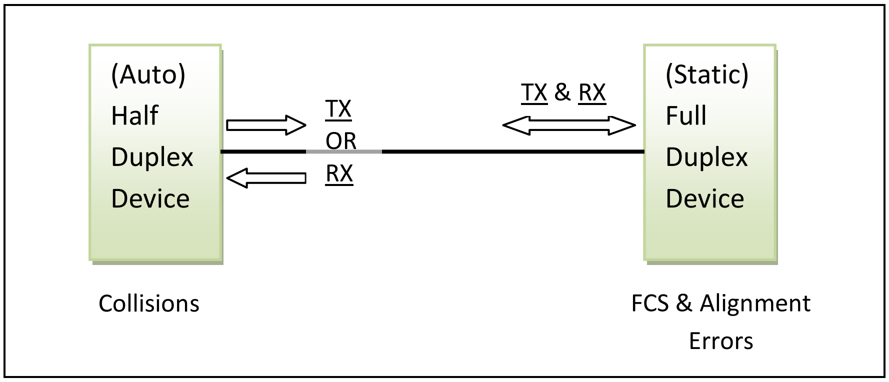

FIGURE 2 - DUPLEX MISMATCH

Most devices implement counters to allow monitoring for these errors. These counters are typically available via either EtherNet/IP or through SNMP. With EtherNet/IP they are found in the Ethernet Link Object (0xF6) as optional attributes. With SNMP these individual counters might be provided if the device supports the “etherLike MIB” {RFC 3635}. Or, if not, they are aggregated to mandatory objects that are available via MIB II {RFC 1213}. Most managed switches will support the etherLike MIB; most end devices will not.

The following table shows the correlation of these different access mechanisms:

EtherNet/IP		**SNMP**	
Ethernet Link Object (0xF6)		**RFC 3635 - etherLike MIB (1.3.6.1.2.1.10.7)**	**RFC 1213 – MIB II (1.3.6.1.2.1)**
Attribute ID:	**Name**		
2	Interface Flags (Duplex configuration and operational status)	Dot3StatsDuplexStatus (1.3.6.1.2.1.10.7.2.1.19)	None[1]
5	Alignment Errors	dot3StatsAlignmentErrors (1.3.6.1.2.1.10.7.2.1.2)	ifInErrors[2] (1.3.6.1.2.1.2.2.1.14)
	FCS Errors	dot3StatsFCSErrors (1.3.6.1.2.1.10.7.2.1.3)	
	Late Collisions	dot3StatsLateCollisions (1.3.6.1.2.1.10.7.2.1.8)	ifOutErrors[3] (1.3.6.1.2.1.2.2.1.20)
	Excessive Collisions	dot3StatsExcessiveCollisions (1.3.6.1.2.1.10.7.2.1.9)	

TABLE 3 - DUPLEX MISMATCH COUNTER CORRELATION

Note 1: Some devices will provide the duplex setting and status information via their private MIB. Check the user documentation on the device.

Note 2: ifInErrors = Alignment Errors + FCS Errors + Frame Too Long Errors + Internal MAC Receive Errors. In this scenario Frame Too Long and Internal MAC Receive errors will be less likely and are ignored.

Note 3: ifOutErrors = Late Collisions + Excessive Collisions + SQE Test Errors + Internal MAC Transmit Errors + Carrier Sense Errors. In this scenario SQE Test Errors, Internal MAC Transmit Errors, and Carrier Sense Errors will be less likely and are ignored.

Some network engineers subscribe to the notion of always using statically configured half duplex links. Then it does not matter as much if a device is misconfigured for AutoNeg as it will revert to half duplex anyway and the link will operate as expected. The performance will be degraded over full duplex but significantly better than a duplex mismatch link.

6.3. Bad Performance Caused by Overloaded Devices

Many industrial network end devices may have a throughput limitation that is significantly less than the connection speed. A device may establish an Ethernet link with a 100 MB/second link speed but this does not imply that the device can maintain a sustained 100 MB/second throughput; likely much less. Often network traffic can become “bursty” with short peaks of high network utilization. The greater the difference between the link speed and the “maximum sustainable rate”, the more susceptible a device may be to these bursts.

The maximum sustainable rate can be very difficult to quantify. It may depend less on the data rate of inbound packets and more on the traffic profile as a collection of mostly small packets, mostly large packets, or a mixture of each. Also the packet types (broadcast, multicast, unicast) and the protocol come into play as this affects the processing time of the received packets.

Most devices use memory buffers to allow them to support a short burst of traffic at the full line rate. However, the device will likely not be able to process the incoming traffic at near that speed thus filling up the buffers. It is not only the rate at which packets arrive but also the size of the packets that matter as it comes down to the available storage in the typically fixed size buffers. Keep in mind that the packets in these buffers are not only unicast traffic that is destined for the device, but also broadcast and potentially multicast traffic as well. Actually, problems with CPU overload due to traffic are most likely caused by multicast and not unicast, as unicast that is not directed at the device MAC/CPU, is dropped at the MAC layer.

If a device's buffers fill up it will be forced to drop otherwise perfectly good packets as there is no more room to store the new packets. When this happens the device should increment a discard counter to indicate that a good packet was thrown away. This counter is available via the EtherNet/IP Ethernet Link Object (0xF6) – Attribute 4 – "In Discards". Alternatively this counter is available through SNMP MIB II – ifInDiscards (1.3.6.1.2.1.2.2.1.13.0). If this counter is increasing the device is overloaded.

The symptoms of this overload may vary but typically the overload can result in sluggish performance. The dropped packets cause delays in the communication stream but not necessarily connection timeouts. If a device is experiencing network traffic at a high enough level to cause discards, its CPU utilization is probably also very high; thus causing the device to be even more sluggish. The CPU utilization is not as readily available via SNMP or EtherNet/IP but it is supported by many devices that have web pages. In EtherNet/IP it is an optional attribute of the Connection Manager Object (0x06) – attribute 11. In SNMP some devices provide it via a private MIB.

If the traffic load becomes very excessive it is possible to drop enough packets that connections would begin timing out. It is important to recognize that connections that have been operating without error for a long time and suddenly become sluggish could be getting affected by external influences such as high broadcast rates or leakage of multicast traffic. The overload condition could be a transient situation caused, for example, by a system startup where multiple connections are starting simultaneously. The overload may also be caused by other situations in the network such as an excessive rate of broadcast traffic or multicast that is not properly filtered.

Multicast traffic in the network can be controlled with the use of IGMP Snooping. IGMP Snooping "prunes" the flow of multicast traffic off of links for which there are no downstream subscribers to the multicast traffic stream. Therefore, ideally, after pruning multicast traffic will not be sent to any devices that are not configured to receive the traffic. This is only true if all infrastructure devices, including multiport devices, support IGMP Snooping and at least one device provides IGMP Querier capability. Also, IGMP Snooping can take minutes to complete the pruning process and during this time the multicast traffic is flooded to all devices. If the problem resolves itself after a few minutes, check the status of IGMP Snooping.

IGMP Snooping must also recalculate the pruning with every topology change of a redundant network such as with RSTP or DLR. If the overload situation is coincident with a redundancy topology change it could be a side effect of some other situation in the network that caused a topology change. Both problems must be addressed separately to ensure the root cause of each is found.

There are methods to help mitigate the external influences to end devices. First and foremost, the use of managed switches with some traffic shaping control will help to shield these devices from the excessive undesirable traffic. Broadcast and multicast rate limiting at an upstream switch will help to curtail this flow of traffic that may be filling up the device's buffers.

6.4. Bad Performance Caused by Electrical Noise

Electrical noise injected into a cable can interfere with the Ethernet signals on the cable and cause packet corruption. While Ethernet cable quality continues to improve, performance problems attributed to induced electrical noise on the wire remains to be a problem in some industrial environments. Often these problems are caused by installation problems leading to failing cables or connectors. In a high electrical noise environment it is crucial to follow every aspect of the cabling guidelines set forth in "EtherNetIP_Media_Planning_and_Installation_Manual".

Injected electrical noise can cause bits within the communication stream to change and corrupt the packet. These bit changes cause FCS errors and Alignment errors in the received packet. When these corrupted packets are detected at the receiving device they are discarded.

Packet corruption can be caused by a malfunctioning device. Often when devices fail they stop communicating all together and sometimes will not even establish an Ethernet link. It does happen occasionally that a device will fail in such a manner that it injects noise or otherwise corrupts packets on the wire. However, a packet corruption problem is more often attributable to a cabling problem.

Cabling is often a blind item and the fault may actually be somewhere hidden to the maintenance personnel. Noise injection may also be a fault of cable routing as, for example, the cabling may be very close to other high conducting cables such as welding leads or main power feeds. Exceedingly tight bends in the cable are also common. These bends in the cable can breakdown insulators within the cable and deform twisted pairs; thus allowing crosstalk between pairs, electrical noise induction or impedance mismatch. Corrosion of the conductors over time can cause the signal to degrade as well. Another common source of noise is improperly applied cable terminations. Sometimes during a connector installation an excessive length of cable is untwisted to make it easier to install the connector. This untwisting of the pairs diminishes the cables natural noise cancelation effect and reduces its noise immunity.

Proper shielding and grounding techniques are vital to robust communications in a high noise environment. The two together, shielding and grounding are critical as shielding alone can cause more problems than it solves. Without proper grounding the ensuing ground loop currents can inject more noise rather than reduce noise.

In situations with high electrical noise or potential ground loop problems, it may be worthwhile to consider the use of fiber optic cabling for the Ethernet communications as it is immune to these disturbances. Also fiber can significantly increase the length of cable runs and improve bandwidth availability. Fiber, however, comes with its own set of installation requirements such as minimum bend radiuses that to prevent signal attenuation and protection from physical shock that can cause it to shatter. Also properly applying cable terminations to fiber requires significantly more skill than twisted pair copper. Again, refer to "EtherNetIP_Media_Planning_and_Installation_Manual" for proper installation requirements.

It can be difficult to distinguish between an electrical noise problem and a faulty cable or connector. The faulty cable can reduce the signal to noise ratio or have other defects that allow noise interference. The faulty cable's noise immunity may be compromised. Using a good quality cable tester can be an invaluable first step in diagnosing suspected noise and cable fault problems.

Electrical noise and cable fault problems can be difficult to detect with Ping or packet capture. Since electrical noise is stochastic, the probability that a particular ICMP (ping) packet will be affected might be slim. Ping may indicate that the network is healthy when in fact there may be many packets being lost. Using packet capture to try to pinpoint the location of the link that is inducing the noise can be futile. As packets traverse a network, each switch along the path will check the packet and if it is corrupted the packet gets discarded at the inbound interface. A packet capture computer connected through a mirrored port will never see the corrupted packets; it will only see the results of the lost packets with no indication why they are lost.

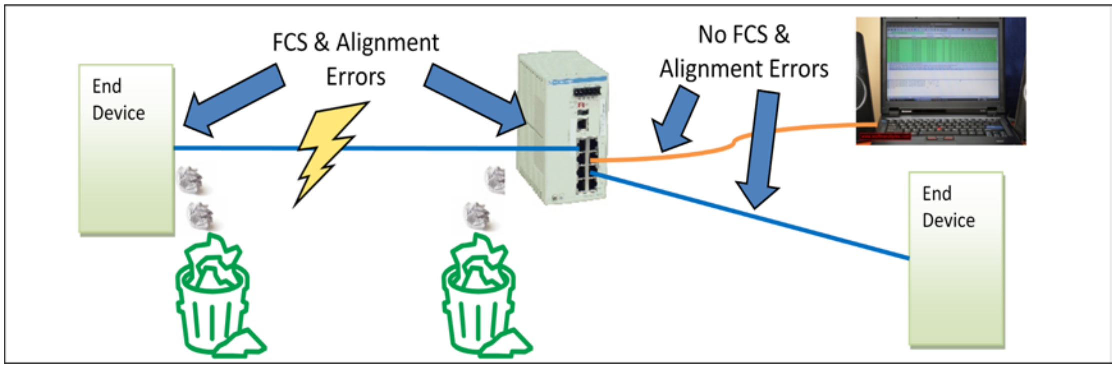

FIGURE 3 - ELECTRICAL NOISE ON A LINK

On a full duplex configured link, the key to distinguishing between cable fault and electrical noise problems from duplex mismatch problems is the absence of collisions. Both classes of problems will cause FCS and Alignment errors but with a duplex mismatch, as discussed previously, there will be collisions.

On a half duplex link, the key to distinguishing noise induced errors from other errors is to note the collision counter values. In a proper half duplex link both ends will detect collisions; this is perfectly normal and expected. And the majority of these collisions should be normal collisions (Single Collisions or Multiple Collisions). But the relative number of FCS and Alignment errors should be much less. If the link is detecting a high rate of FCS or alignment errors relative to the number of collisions, suspect an electrical noise problem.

In some cases it might be possible to infer the nature of the problem by a closer examination of the relative counter values. Consider the case of a bad cable where only one pair is damaged. Maybe the cable was crushed or bent to sharply. In this scenario, since each pair is used for a single direction, the errors may be more prevalent on one end of the link. If this were a noise injection problem then it likely would not affect only a single pair. If these errors are primarily detected on only one end, suspect a faulty cable.

Both EtherNet/IP Ethernet Link Object (0xF6) and SNMP provide access to the Ethernet counters. See Table 5 - EtherNet/IP to SNMP Correlation. The key counters to look for are the FCS and Alignment errors along with a comparison of the relative value of the collision counters. Remember, it is the rate that matters, not the raw counter value.

6.5. Bad Performance Caused by Excessive Cable Length

Excessive cable length is another possible cause of bad performance. Although the failure modes and indicators are a little different in half duplex versus full duplex, the maximum standard cable length is the same for twisted pair copper. Both mediums, half and full duplex links, will suffer from increased attenuation causing reduced signal strength in an excessively long cable. This reduces the signal to noise ratio (SNR) and makes the signal much more susceptible to noise interference. This interference causes FCS and Alignment errors in packets; thus these packets are dropped. Statistically, these errors will be agnostic to the direction of data so will appear on both ends of the link.

Half duplex has an additional failure mode where packet propagation delays cause collisions to be detected after the 64 byte window known as the slot time. Collisions that occur after the slot time are detected as late collisions and are dropped. In a half duplex link if there are a significant number of late collisions then suspect that the cable is longer than the IEEE 802.3 specified length.

6.6. Impact of Hubs in the Network

Ethernet hubs, as opposed to switches, are an antiquated technology. Hubs are by definition a shared media device and with the use of shared media in Ethernet come a number of performance degrading behaviors. Shared media devices require the use of half duplex communications. And with half duplex comes the problem of collisions that, in the best case, cause back-off and retransmission delays and in the worst case, packet loss. Every device in a hub connected network will see every other devices traffic. And every device must contend with this traffic every time it needs to transmit a packet on the wire. Collisions abound and performance suffers. For this reason hubs have been virtually eradicated in enterprise level networks. Very few vendors still sell hubs today.

Industrial Ethernet networks can be much slower to evolve as the “if it isn’t broke, don’t fix it” mentality prevails. But as the industrial Ethernet network continues to grow, so does the traffic. As traffic grows so does the probability of collisions and communication performance declines accordingly. The lesser level of traffic that used to work fine through these hubs may have evolved into levels that incur a high rate of performance robbing collisions. Personal experience has shown that in a moderately loaded link as much as a 30% performance improvement can be gained by changing from hubs to switches at the same speed.

Most hubs are unmanaged devices that will offer no access to diagnostic data. Therefore, collecting remote diagnostic information from hub to hub links is not possible. But if collisions are happening at one end of a link then they are likely also happening at the other end of the link. Therefore, hub to device links can be analyzed for collisions by viewing information from only the device end of the link. This can be done with either EtherNet/IP or SNMP.

Because of the shared media technology, the use of “temporary” hubs can offer a quick and dirty way to perform packet capture. By injecting a hub into the middle of a link and then connecting an analysis machine, packet capture of all traffic on that link is possible. It is very important to recognize that by doing so the system has been changed; it is not the same system that originally exhibited the problem to be investigated. Especially if this link was a full duplex link, it must now contend with collisions in addition to whatever else was happening on the link that warranted the investigation. This could change the behavior to a degree that leads to an incorrect conclusion from the analysis.

6.7. Profiling the Different Types of Errors

In profiling the different types of errors it can be difficult to apply absolute limit values to the counter rates. The rates that indicate problems are more of a generalized rule of thumb. As a general rule, any collision rate higher than 1% of network traffic is suspect. It is important to remember that normal collisions are not as detrimental as any of the error collisions, FCS and Alignment errors, and Discards. With normal collisions the packets are still getting through the network, with any of the other errors the packets are lost.

For duplex mismatch there will likely be more normal collisions than error collisions on the AutoNeg end of the link. And more FCS and Alignment errors on the Static end than on the AutoNeg end. There should be few Discards as the link will not be performing well enough to cause an overload of the device.

Overloaded devices with half duplex links will likely have some collisions as the link is busy; how busy depends partly on the throughput of the device. Of course full duplex links will have zero collisions. The important indicator is that the Discards rate is non-zero.

Electrical noise and cable faults on half duplex links will have slightly elevated levels of both normal and error collisions as the noise can be misinterpreted as a collision. Again, full duplex links will have zero collisions. But much higher levels of FCS and Alignment Errors are the important indicators.

Long cables cause problems similar to noise problems except that, for half duplex links, they will likely have an elevated level of late collisions as the propagation delay causes the collisions to be detected after the slot time.

The values given in the table below (see Table 4 - counter value Profiles) are generalized and relative to normal operation and the other counters considered for each case.

- Low indicates normal operation
- Med indicates higher than normal – maybe 2%
- High indicates anything above Med

Problem Type		**Normal Collisions Rate**	**Error Collisions Rate**	**FCS / Alignment Errors Rate**	**Discards Rate**
Duplex Mismatch	Static Config (FD) End	0	0	High	Zero to few
	AutoNeg (HD) End	High	Med	Low	
Overloaded Devices	Half Duplex	Med to High	Med to High	Low	Rate >> 0 indicates potential problem
	Full Duplex	0	0		
Noise / Cable Fault	Half Duplex	Med	Med	High	Zero to few
	Full Duplex	0	0		
Cable too long	Half Duplex	Med	High	High	Zero to few
	Full Duplex	0	0	High	

TABLE 4 - COUNTER VALUE PROFILES

7. Appendix

7.1. EtherNet/IP Diagnostics and SNMP Correlation

EtherNet/IP Ethernet Link Object (0xF6)			SNMP	
Attribute ID:	**Name**		**RFC, MIB, & Table**	**MIB Object & (OID)**
1	Interface Speed		RFC 1213 MIB II Interface Table - ifTable (1.3.6.1.2.1.2.2.1.x)	ifSpeed (ifTable.5)
2	Interface Flags			None
3	Physical Address			ifPhysAddress (ifTable.6)
4	Interface Counters	In Octets		ifInOctets (ifTable.10)
		In Ucast Packets		ifInUcastPkts (ifTable.11)
		In NUcast Packets		ifInNUcastPkts (ifTable.12)
		In Discards		ifInDiscards (ifTable.13)
		In Errors		ifInErrors (ifTable.14)
		In Unkown Protos		ifInUnknownProtos (ifTable.15)
		Out Octets		ifOutOctets (ifTable.16)
		Out Ucast Packets		ifOutUcastPkts (ifTable.17)
		Out NUcast Packets		ifOutNUcastPkts (ifTable.18)
		Out Discards		ifOutDiscards (ifTable.19)
		Out Errors		ifOutErrors (ifTable.20)

EtherNet/IP Ethernet Link Object (0xF6)			SNMP	
Attribute ID:	Name		RFC, MIB, & Table	MIB Object & (OID)
5	Media Counters (Based on RFC 1643)	Alignment Errors	RFC 3635 etherLike MIB dot3StatsTable (1.3.6.1.2.1.10.7.2.1.x)	dot3StatsAlignmentErrors (dot3StatsTable.2)
		FCS Errors		dot3StatsFCSErrors (dot3StatsTable.3)
		Single Collisions		dot3StatsSingleCollisionFrames (dot3StatsTable.4)
		Multiple Collisions		dot3StatsMultipleCollisionFrames (dot3StatsTable.5)
		SQE Test Errors		dot3StatsSQETestErrors (dot3StatsTable.6)
		Deferred Transmissions		dot3StatsDeferredTransmissions (dot3StatsTable.7)
		Late Collisions		dot3StatsLateCollisions (dot3StatsTable.8)
		Excessive Collisions		dot3StatsExcessiveCollisions (dot3StatsTable.9)
		MAC Transmit Errors		dot3StatsInternalMacTransmitErrors (dot3StatsTable.10)
		Carrier Sense Errors		dot3StatsCarrierSenseErrors (dot3StatsTable.11)
		Frame Too Long		dot3StatsFrameTooLongs (dot3StatsTable.13)
		MAC Receive Errors		dot3StatsInternalMacReceiveErrors (dot3StatsTable.16)

TABLE 5 - ETHERNET/IP TO SNMP CORRELATION

8. References

The CIP Networks Library Volume 2 EtherNet/IP Adaption of CIP – Version 1.15 / April 2013

PUB00148R0_EtherNetIP_Media_Planning_and_Installation_Manual.pdf

RFC 1213 MIB-II Management Information Base

RFC 1643 - Definitions of Managed Objects for the Ethernet-like Interface Types

RFC 3635 - Definitions of Managed Objects for the Ethernet-like Interface Types

RFC 3638 - Applicability Statement for Reclassification of RFC 1643 to Historic Status

http://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr1904.html

http://www.brocade.com/downloads/documents/html_product_manuals/B6910_DIAG_2104/wwhelp/wwhimpl/common/html/wwhelp.htm#href=L1_diagnostics.5.2.html&single=true

http://www.techfest.com/networking/lan/ethernet3.htm

https://www.appliedtrust.com/resources/performance/untangling-ethernet-performance-problems

https://supportforums.cisco.com/docs/DOC-1806

****

The ideas, opinions, and recommendations expressed herein are intended to describe concepts of the author(s) for the possible use of CIP Networks and do not reflect the ideas, opinions, and recommendation of ODVA per se. Because CIP Networks may be applied in many diverse situations and in conjunction with products and systems from multiple vendors, the reader and those responsible for specifying CIP Networks must determine for themselves the suitability and the suitability of ideas, opinions, and recommendations expressed herein for intended use.