Network Infrastructure for EtherNet/IP™: Introduction and Considerations

Network Infrastructure for EtherNet/IP™
Publication Number: PUB00035R0

ODVA, Inc.
4220 Varsity Drive, Suite A
Ann Arbor, MI 48108-5006 USA
Phone +1 734.975.8840
Fax +1 734.922.0027
E-mail odva@odva.org
Web www.odva.org

The right to make, use, or sell product or system implementations described herein is granted only under separate license pursuant to a Terms of Usage Agreement or other agreement. Terms of Usage Agreements for individual CIP Networks are available, at standard charges, over the Internet at the following web sites:

www.odva.org Terms of Usage Agreements for DeviceNet, CompoNet, and EtherNet/IP along with general information on DeviceNet, CompoNet, and EtherNet/IP networks and their extensions, and the association of ODVA.

www.controlnet.org Terms of Usage Agreement for ControlNet along with general information on ControlNet and ControlNet International.

NOTE: Because the technologies described herein may be applied in many diverse situations and in conjunction with products and systems from multiple vendors, the user and those responsible for specifying these technologies must determine for themselves their suitability for the intended use. ALL INFORMATION PROVIDED BY ODVA IS PROVIDED ON AN “AS IS” BASIS. NO WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING WITHOUT LIMITATION ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, ARE BEING PROVIDED BY THE PUBLISHER OR ODVA. The figures and examples in this guide are designed to demonstrate general concepts for planning and installing a network. The user should always verify interconnection requirements to and from other equipment, and confirm installation and maintenance requirements for their specific application. IN NO EVENT SHALL THE PUBLISHER OR ODVA, THEIR OFFICERS, DIRECTORS, MEMBERS, AGENTS, LICENSORS, OR AFFILIATES BE LIABLE TO YOU, ANY CUSTOMER, OR THIRD PARTY FOR ANY DAMAGES, DIRECT OR INDIRECT, INCLUDING BUT NOT LIMITED TO LOST PROFITS, DEVELOPMENT EXPENSES, OR ANY OTHER DIRECT, INDIRECT, INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES.

Note from the Publisher

The application of Ethernet and Internet technologies is permeating every aspect of manufacturing automation today—control, safety, configuration and diagnostics, synchronization and motion, and information. To realize the benefits of these technologies, the network infrastructure is critical and represents a long-term investment for users that requires thorough analysis of performance requirements along with technology choices and trends.

Network Infrastructure for EtherNet/IP™ provides an introduction to the network infrastructure used in EtherNet/IP networks and gives users a framework for identifying the considerations that are most critical to their specific applications. It represents the combined knowledge of experts from companies that are leading the industry in the application of Ethernet on the factory floor. ODVA wishes to recognize the many hours of hard work provided by people from these contributing companies:

- Belden
- Cisco Systems
- Contemporary Controls
- General Motors
- HARTING
- Hirschmann
- Kendall Electric
- McNaughton-McKay Electric Company
- Phoenix Contact
- Rockwell Automation
- RuggedCom
- The Siemon Company
- WAGO

To learn more about EtherNet/IP, visit the ODVA web site at www.odva.org where you'll find many more resources to help you apply the world's leading Ethernet network for manufacturing automation.

Sincerely,

Katherine Voss

Katherine Voss
Executive Director
ODVA

Contents

Preface

About This Publication

Advances in Ethernet technology and reduced component costs are fueling the growth of Ethernet in plant floor networks. EtherNet/IP™, the adaptation of the Common Industrial Protocol™ (CIP™) provides users with the tools to deploy standard Ethernet technology for manufacturing applications while enabling Internet and enterprise connectivity… anytime, anywhere. The unique performance and environmental requirements of manufacturing applications require users to select appropriate network infrastructure. This publication provides an overview of the technologies used to deploy EtherNet/IP and provides guidelines for deploying infrastructure devices in EtherNet/IP networks.

Related Publications

For more information on how to plan for, design, implement, and maintain networks incorporating EtherNet/IP, refer to the current versions of the following publications:

- ODVA, *The EtherNet/IP™ Specification*
- ODVA, *EtherNet/IP™ Media Planning and Installation Guide*
- IAONA, *Industrial Ethernet Planning and Installation Guide*
- IEEE 802®, IEEE Standard for Information Technology –Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Networks

1 Industrial Ethernet... Not Just Another Fieldbus

Using the Common Industrial Protocol (CIP™), EtherNet/IP provides the backbone for a completely open, vendor-neutral communication network for the full range of manufacturing applications, and can integrate these applications with one another and the enterprise. Unlike traditional fieldbuses, EtherNet/IP has an active infrastructure that uses the standard Ethernet physical layer and standard Ethernet and Internet protocols. This chapter provides the background to understand the application of industrial networks and their special needs along with an introduction to the unique aspects of deploying Ethernet networks.

1.1 Industrial Networks

Industrial communication networks have been deployed on the factory floor since the 1980s. Initially used as remote input/output (I/O) and controller-level networks, they now interconnect a wide range of industrial automation devices, from sensors and actuators, to controllers, human-machine interface (HMI) workstations, and computers. Today, these networks are capable of supporting real-time control as well as plant-floor information systems.

The characteristics of industrial networks differ from commercial or residential networks in two ways:

1. They provide the additional real-time performance capabilities needed for a majority of manufacturing applications.

2. They meet the requirements for survival in various types of industrial environments, normally characterized by higher levels of electrical noise, shock, vibrations, ambient temperature, humidity, etc.

1.2 Introduction to EtherNet/IP

By the early 1990s, Ethernet networks were used successfully in manufacturing environments, mainly in non-time-critical applications. The evolution of the Ethernet technology from a 10-Mbps, half-duplex, bus/tree topology to a 100-Mbps and 1-Gbps, full-duplex, switch/router-based star topology has paved the way for using Ethernet to support time-critical applications in industrial networks.

EtherNet/IP is a communication network designed for use in industrial environments and time-critical applications. Using the Common Industrial Protocol™ (CIP™), it has many unique features that give it advantages over other Ethernet-based industrial networks.

EtherNet/IP:

- offers producer-consumer services that enable users to control, configure, and collect data from over a single network or to use the network as a backbone for multiple distributed networks.

- provides robust physical layer options for industrial environments and includes the use of sealed RJ-45 and M12-4 "D" Coded connectors.

- uses existing Institute of Electrical and Electronics Engineers (IEEE) standards for Ethernet physical and data link layers.

- provides flexible installation options leveraging commercially available industrial infrastructure products, including copper, fiber, fiber ring, and wireless solutions.

- incorporates the TCP/IP suite—the Ethernet standard.

- is compatible with accepted communication standards, including OPC[1] and Internet protocols such as Hypertext Transfer Protocol (HTTP), Simple Network Management Protocol (SNMP), File Transfer Protocol (FTP), and Dynamic Host Configuration Protocol (DHCP).

It is useful to understand the architecture of EtherNet/IP in terms of the seven-layer Open System Interconnection (OSI) Reference Model developed by the International Organization for Standardization (ISO) as shown in Figure 1-1. *Appendix B* provides an introduction to the ISO OSI Model with references to appropriate standards.

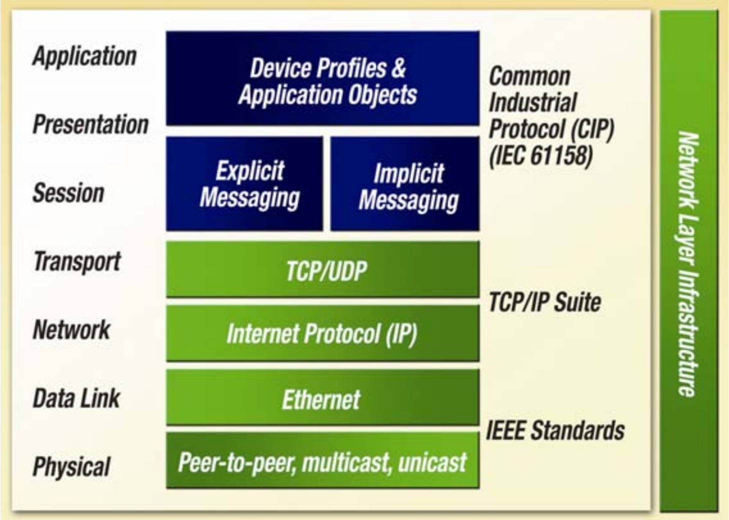

Figure 1-1 How EtherNet/IP Relates to the OSI Model.

In EtherNet/IP and all CIP Networks, exchange of time-critical data is based on the producer-consumer messaging model whereby a transmitting device produces data on the network, and one or more

[1] Information on the OPC Foundation and its open specifications can be found at www.opcfoundation.org.

receiving devices can consume this data simultaneously. Implementation of this model in CIP is supported by the Internet Protocol (IP) multicast service which, in its turn, is supported by the Ethernet multicast service. The primary benefit of a producer-consumer network is its more efficient use of bandwidth. When a message is produced on the network, it is identified not by its destination address, but by its connection ID. Multiple nodes may then consume the data to which the connection ID refers. As a result, if a node wants to receive data, it only needs to ask for it once in order to consume the data each time it is produced. And, if a second (third, fourth, etc.) node wants the same data, all it needs to know is the connection ID to receive the same data simultaneously with all other nodes.

Conversely, using the source-destination model, nodes receive only the packets that contain their destination node number. If more than one node needs the same data, it must be transmitted multiple times, which is inherently inefficient. This can also cause synchronization problems, as nodes that require the same data obtain it at slightly different times.

1.2.1 Industrial Ethernet and Traditional Fieldbuses

Since the mid-1990s, the use of fieldbuses, such as DeviceNet™, has grown significantly. At the same time, advances in Ethernet switching technology and significant cost reductions fueled by the expanded use of PCs and the Internet now allow industrial Ethernet to be used in more manufacturing applications. Depending on the application, it may be desirable to combine Ethernet networks with fieldbus-type networks. CIP Networks, such as EtherNet/IP and DeviceNet, support seamless bridging and routing between multiple CIP Networks.

Implementing an industrial Ethernet system is different than implementing a device-level network, such as DeviceNet. Whether installing a homogeneous EtherNet/IP network architecture or a heterogeneous CIP Network architecture, users must be aware of the basic differences between Ethernet networks, such as EtherNet/IP, and typical fieldbuses, such as DeviceNet, as highlighted in Table 1-1.

Table 1-1 Comparison of Device-level Networks and Industrial Ethernet Implementations

Control System Issue	Typical Device-level Network Capability	Industrial Ethernet Capability	Industrial Ethernet Differences
Device capacity and wiring flexibility	Trunk-and-drop, typically 10s of devices Limited total network distance	Star configuration with potential for 100s or 1000s of devices Virtually unlimited total network distance	Infrastructure devices required Network management tools may be needed
Data rate vs. distance	100s kbps, up to 500 meters (1600 ft) for copper cable runs, trading lower speed for longer distance	10 Mbps to 100 Mbps, up to 100 meters (328 ft) for copper runs and up to 2000 meters (6561 ft) for fiber	Cable types and noise immunity Greater need for a mix of fiber and copper within a system
Protocols	Single protocol	Coexistence of multiple protocols	Coexistence of multiple protocols

1.2.2 Infrastructure for EtherNet/IP Applications

A communication network can be viewed as a utility for providing services that allow devices, or humans operating these devices, to exchange information. A typical modern communication network is functionally a mix of computer, telecommunication, and broadcast networks capable of supporting voice, data, video, and image communication. Such a network normally consists of two basic components: nodes and infrastructure (i.e., network = nodes + infrastructure). The same is true for industrial networks.

In the world of industrial networks, two types of infrastructures can be identified: passive and active. Fieldbuses typically have a passive infrastructure, which is essentially a cable system. In contrast, Ethernet networks, including those using EtherNet/IP, have an active infrastructure. At its core is an interconnection of Ethernet Layer 2 (OSI data link layer) and Layer 3 (OSI network layer) switches. Layer 2 switches allow all protocols to pass, but Layer 3 switches behave like routers, allowing only designated protocols to pass through to certain locations.

The following examples present approaches that could be used to design a network with active infrastructure. For Ethernet to function in control applications, it has to satisfy end-to-end response time requirements for each transaction. Because processing delays in the network infrastructure are usually negligible in comparison with delays in end-devices (nodes), performance bottlenecks are still located in end-devices. Some of these end-devices (e.g., controllers) have to support time-critical as well as non-time-critical traffic and are frequently equipped with two separate Ethernet interfaces. But there is no need to create two physical networks to separate devices with time-critical data from devices with non-time-critical traffic. This can be achieved by applying switches that support virtual LAN (VLAN) functionality, thus making it possible to create separate virtual networks on the same wire. The performance of many commercial and industrial switches manufactured today is sufficient to handle this combined traffic.

When an EtherNet/IP network is integrated with a plant’s enterprise network, network designers must consider a set of issues never previously considered by designers of industrial networks. For instance, EtherNet/IP infrastructure devices must have features that, when enabled, will not allow propagation of time-critical traffic onto the plant’s enterprise network. This would bog down the enterprise network. Conversely, extraneous enterprise network traffic must not propagate onto the EtherNet/IP network.

1.2.3 Security Issues

The deployment of Ethernet networks for manufacturing applications brings all of the issues and considerations for network security to the factory floor. The coexistence of multiple protocols and IP addressing mean that network systems designers must take additional steps to ensure secure operation of the manufacturing applications. Not only is it important that enterprise traffic be kept out of the plant-floor network, but it is also important to keep viruses, Trojans, and worms from infecting manufacturing computers. In addition, access to the plant-floor network controllers must be restricted to prevent inappropriate changes in controller configurations and programs.

Fortunately, the active infrastructure needed for EtherNet/IP can help factory-floor applications be more secure. Proper installation of managed (configurable) switches, routers, and firewalls provide mechanisms for managing traffic to achieve a more secure network installation.

2 Understanding the Basics of Network Protocols

The following discussion includes information on the TCP/IP suite, focusing on the Transmission Control Protocol (TCP), the Internet Protocol (IP) and the User Datagram Protocol (UDP).

The OSI Model provides a conceptual framework for communication between computers, but the model itself is not a method of communication. Actual communication is made possible by using communication protocols. In the context of data networking, a protocol is a formal set of rules and conventions that govern how computers exchange information over a network medium. A protocol implements the functions of one or more of the OSI Model layers. (See *Appendix B* for more information.)

A wide variety of communication protocols exist. Some of these protocols include Local Area Network (LAN) protocols, Wide Area Network (WAN) protocols, network protocols, and routing protocols. LAN and WAN protocols operate at the physical layer (Layer 1) and data link layer (Layer 2) of the OSI Model and define communication over the various LAN media. Routing protocols are network-layer protocols that are responsible for exchanging information between routers, so the routers can select the proper path for network traffic. Finally, network protocols are the various upper-layer protocols that exist in a given protocol suite. (See Figure 2-1.)

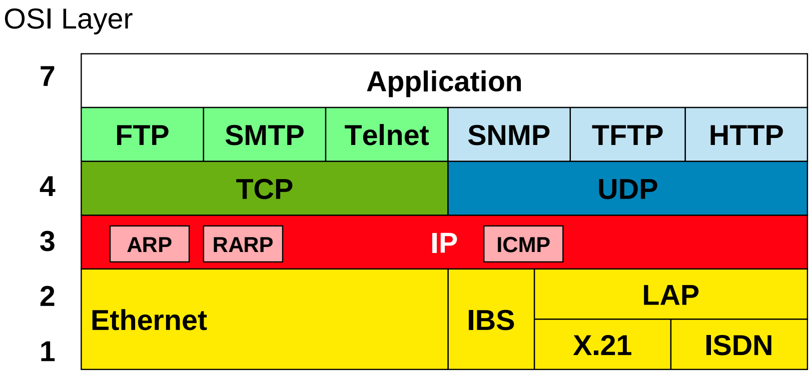

Figure 2-1 Relationship of the Internet Protocol Suite to the OSI Reference Model.

Many protocols rely on the existence of others for operation. For example, CIP uses Internet Protocol (IP), UDP, and TCP. This concept of building upon other existing layers is the foundation of the OSI Model.

2.1 The TCP/IP Suite

EtherNet/IP incorporates the TCP/IP suite, utilizing TCP for explicit messaging and UDP for implicit or I/O messaging. These widely adopted protocols are implemented at the network layer (IP) and the transport layer (UDP and TCP) of the OSI Model, and are the basic networking technology for EtherNet/IP.

Transmission Control Protocol (TCP)

Transmission Control Protocol (TCP) is a connection-oriented transport protocol that sends data as an unstructured stream of bytes. By using sequence numbers and acknowledgment messages, TCP can provide a sending node with delivery information about packets transmitted to a destination node. EtherNet/IP uses TCP for explicit messaging (i.e., device configuration).

TCP's multiplexing allows simultaneous upper-layer conversations to be multiplexed over a single connection. TCP provides connection-oriented, end-to-end reliable packet delivery by sequencing bytes with a forwarding acknowledgment number that indicates to the destination the next byte the source expects to receive. Bytes not acknowledged within a specified time period are retransmitted. The reliability mechanism of TCP allows devices to deal with lost, delayed, duplicate, or misread packets.

For example, if data is lost in transit from source to destination, TCP can retransmit the data until either a timeout condition is reached or successful delivery is achieved. TCP can also recognize duplicate messages and appropriately discard them. If the sending computer is transmitting too fast for the receiving computer, TCP can employ flow control mechanisms to slow data transfer. TCP can also communicate delivery information to upper-layer protocols and applications. Figure 2-2 shows the TCP packet fields.

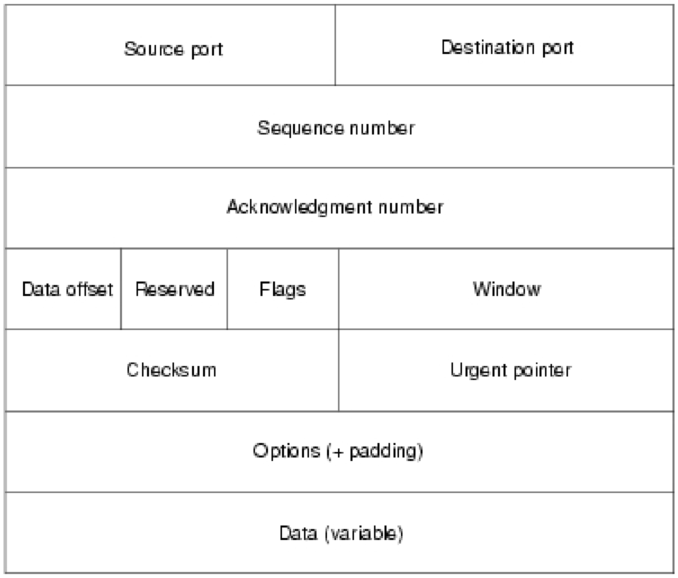

Figure 2-2 TCP Packet Fields.

The following descriptions summarize the TCP packet fields.

- Source Port and Destination Port: Identifies the source and destination processes that are to receive TCP services.
- Sequence Number: Typically specifies the number assigned to the first byte of data in the current message.
- Acknowledgment Number: Contains the sequence number of the next byte of data the sender expects to receive.
- Data Offset: Indicates the number of 32-bit words in the TCP header.
- Reserved: Reserved for future use.
- Flags: Carries a variety of control information used for connection establishment, and termination.
- Window: Specifies the size of the buffer space available for incoming data.
- Checksum: Indicates if a header or data was damaged in transit.
- Urgent Pointer: Points to the first non-urgent data byte in the packet.
- Options: Specifies various TCP options.
- Data: Contains upper-layer information.

Internet Protocol (IP)

Internet Protocol (IP) is the primary Layer 3 protocol in the TCP/IP suite. In addition to inter-network routing, IP provides error reporting as well as fragmentation and reassembly of information units called datagrams for transmission over networks with different maximum data unit sizes. IP represents the foundation of the TCP/IP suite.

Public IP addresses are globally unique, 32-bit numbers assigned by the Network Information Center (NIC). Globally unique addresses permit IP networks anywhere in the world to communicate with each other. In control networks, users may assign IP addresses within “private” IP address ranges. Examples of “private” IP address ranges include 10.x.x.x, 172.16-31.x.x, and 192.168.x.x. The IP addresses in these “private” control networks must be assigned so that they are also unique on the network. The setting of identical addresses on the same network will disrupt proper operation.

An IP address is divided into three parts. The first part designates the network address, the second part designates the subnet address, and the third part designates the host address.

IP addressing supports three different network classes. Class A networks are intended mainly for use with a few very large networks because they provide only 8 bits for the network address field. Class B networks allocate 16 bits, and Class C networks allocate 24 bits for the network address field. Class C networks provide only 8 bits for the host field, so the number of hosts per network may be a limiting factor. In all three cases, the leftmost bit(s) indicate the network class. IP addresses are written in dotted decimal format (e.g., 34.0.0.1.). Table 2-1 shows the address formats for Class A, B, and C IP networks.

Table 2-1 IP Address Formats for Class A, B, and C Networks

Class	Leading Bits	Start	End
Class A	0	0.0.0.0	127.255.255.255
Class B	10	128.0.0.0	191.255.255.255
Class C	110	192.0.0.0	223.255.255.255
Class D (multicast[2])	1110	224.0.0.0	239.255.255.255
Class E (reserved)	1111	240.0.0.0	255.255.255.255

IP networks can be divided into smaller units called subnetworks or subnets. Subnets provide extra flexibility for the network administrator. For example, assume that a network has been assigned a Class A address, and all the nodes on the network use a Class A address. Further assume that the dotted decimal representation of this network's address is 34.0.0.0. (All zeros in the host field of an address specify the entire network.) The administrator can subdivide the network using subnetting. This is done by "borrowing" bits from the host portion of the address and using them as a subnet field.

If the network administrator has chosen to use 8 bits of subnetting, the second octet of a Class A IP address provides the subnet number. In our example, address 34.1.0.0 refers to network 34, subnet 1; address 34.2.0.0 refers to network 34, subnet 2, and so on as shown in Figure 2-3.

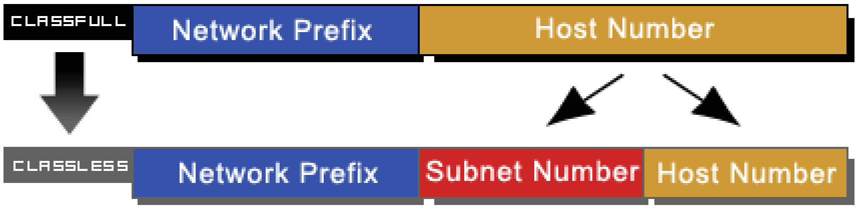

Figure 2-3 Subdividing a Network.

The number of bits that can be borrowed for the subnet address varies. To specify how many bits are used and where they are located in the host field, IP provides subnet masks. Subnet masks use the same format and representation technique as IP addresses as shown in Table 2-2.

[2] A packet with a special destination address, which multiple nodes on the network may be willing to receive. Source: RFC 1983. *Internet Users' Glossary*. G. Malkin. August 1996. http://rfc.net/rfc1983.html

Table 2-2 Subnet Masks

	Dot-decimal Address	Binary
Full Network Address	192.168.5.10	11000000.10101000.00000101.00001010
Subnet Mask	255.255.255.0	11111111.11111111.11111111.00000000
Network Portion	192.168.5.0	11000000.10101000.00000101.00000000

Subnet masks have ones in all bits except those that specify the host field. For example, the subnet mask that specifies 8 bits of subnetting for Class A address 34.0.0.0 is 255.255.0.0. The subnet mask that specifies 16 bits of subnetting for Class A address 34.0.0.0 is 255.255.255.0. Both of these subnet masks are pictured in Figure 2-4. Subnet masks can be passed through a network on demand such that new nodes can learn how many bits of subnetting are being used on their network.

As IP subnets have grown, administrators have looked for ways to use their address space more efficiently. Traditionally, all subnets of the same network number used the same subnet mask. In other words, a network manager would choose an eight-bit mask for all subnets in the network. This strategy is easy to manage for both network administrators and routing protocols. However, this practice wastes address space in some networks.

Some subnets have many hosts and some have only a few, but each consumes an entire subnet number. Serial lines are the most extreme example because each has only two hosts that can be connected via a serial line subnet.

On some media, such as LANs designed to IEEE Std 802®,[3] IP addresses are dynamically discovered through the use of two other members of the Internet protocol suite: Address Resolution Protocol (ARP) and Reverse Address Resolution Protocol (RARP).

ARP uses broadcast messages to determine the hardware layer Media Access Control (MAC) address corresponding to a particular network-layer address. ARP is sufficiently generic to allow use of IP with virtually any type of underlying media-access mechanism.

RARP uses broadcast messages to determine the network-layer address associated with a particular hardware address. RARP is especially important to diskless nodes, for which network-layer addresses usually are unknown at boot time.[4]

2.2 User Datagram Protocol

The User Datagram Protocol (UDP) is a connectionless transport-layer protocol (Layer 4) that belongs to the TCP/IP suite. UDP is an interface between IP and upper-layer processes. UDP protocol ports distinguish multiple applications running on a single device from one another.

[3] IEEE 802, IEEE Standard for Information Technology –Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Networks. http://standards.ieee.org/getieee802/portfolio.html

[4] Information on ARP/RARP courtesy of Cisco Systems, Inc. *Protocol Brief: TCP/IP*. 1996. http://www.cisco.com/warp/public/535/4.html

Unlike the TCP, UDP adds no reliability, flow-control, or error-recovery functions to IP. Because of UDP's simplicity, UDP headers contain fewer bytes and, therefore, consume less network overhead than TCP.

UDP is useful in applications where efficiency of data delivery takes precedence over verification of data delivery, such as in cases where a higher-layer protocol might provide error and flow control.

User Datagram Protocol (UDP)

User Datagram Protocol (UDP) is the transport protocol for several well-known application-layer protocols, including Network File System (NFS), Simple Network Management Protocol (SNMP), Domain Name System (DNS), and Trivial File Transfer Protocol (TFTP). EtherNet/IP uses UDP for implicit (real-time control) messaging.

The UDP packet format contains four fields, as shown in Figure 2-4. These include Source Port, Destination Port, Length, and Checksum fields.

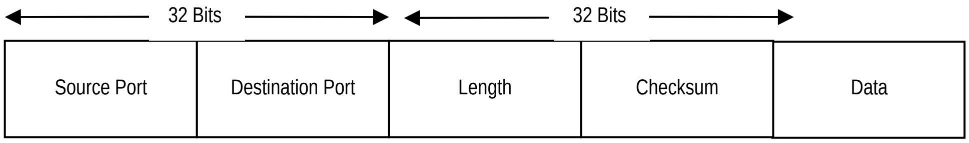

Figure 2-4 UDP Packet Format.

Source Port and Destination Port fields contain the 16-bit UDP protocol port numbers used to de-multiplex datagrams for receiving application-layer processes. A Length field specifies the length of the UDP header and data. Checksum provides an (optional) integrity check on the UDP header and data.

2.3 Data Exchange Models

EtherNet/IP uses the producer-consumer data exchange model. A data exchange model, or architecture, describes a set of mechanisms by which data is exchanged between application programs running in two or more computing devices. Each mechanism, depending on its function, uses unicast, multicast, or broadcast communication or some combination of these. The most popular data exchange models today are point-to-point, client-server, publisher-subscriber, and producer-consumer.

Since present networks support a variety of multimedia information services, it is sometimes difficult to determine whether a particular device supports more than one data exchange model or if one model is a special case of the other one. EtherNet/IP has the flexibility to support these different types of data exchanges to meet application needs that may arise.

Point-to-point communication is the simplest data exchange model. Classical examples of point-to-point connections are communication between two people over a phone and two computers over a TCP/IP connection.

Client-server communication implies a central server connected to many clients. This is a one-to-many model. One could argue that a client-server system is actually a collection of point-to-point connections. Client-server communication (or client-server computing) is popular in office networks and transaction processing systems. Office network examples are file, print, mail, and database servers. Examples related to EtherNet/IP networks are DHCP/BOOTP, DNS, and web servers.

Publish-subscribe communication is a one-to-many model. In these systems, devices subscribe to data they need and publish information they produce. Common examples of publish-subscribe systems include newspapers, magazines, cable television, and various multimedia information services available over the Internet, such as breaking news. The publish-subscribe model is used in some control networks for periodic broadcasts of relatively large amounts of data (e.g., 1000 bytes).

Producer-consumer communication also is a one-to-many model. Information produced by one device can be consumed by a group of other devices. In comparison with the models mentioned above, the producer-consumer model is the most applicable for I/O messaging (transport classes 0 and 1) on EtherNet/IP, which includes cyclic and change-of-state event-type messages, message size, and rate. Refer to *The EtherNet/IP™ Specification*[5] for more details.

Advantages of the Producer-Consumer Model

The primary benefit of a producer-consumer network is its more efficient use of bandwidth. When a message is produced onto the network, it is identified not by its destination address, but by its connection ID. Multiple nodes may then consume the data to which the connection ID refers. As a result, if a node wants to receive data, it only needs to ask for it once in order to consume the data each time it is produced. And, if a second (third, fourth, etc.) node wants the same data, all it needs to know is the connection ID to receive the same data simultaneously with all other nodes. Conversely, using the source-destination model, nodes receive only the packets that contain their destination node number. If more than one node needs the same data, it must be transmitted multiple times, which is inherently inefficient. This can also cause synchronization problems as nodes that require the same data obtain it at slightly different times.

[5] ODVA, The CIP Networks Library, Vol. 1 and 2, *The EtherNet/IP™ Specification.* Ann Arbor: ODVA, Inc., 2006. www.odva.org. CD-ROM.

3 Planning the Infrastructure

The most significant differences between industrial Ethernet networks and device-level networks are presented here with information on how industrial EtherNet/IP networks are designed and built.

Open device-level networks have been a key element of building modular manufacturing systems. The networks' quick response times, low connect costs, and simplicity fuel their growth. At the same time, industrial Ethernet infrastructure components, low-cost embedded Ethernet devices, and better performance due to higher data rates (100 Mbps to 1 Gbps), are increasing the real-time control capability of industrial Ethernet.

3.1 Segmentation vs. Trunk-and-Drop

There are significant differences in the topology of industrial Ethernet networks and device-level networks as shown in Figure 3-1. Device-level networks are often laid out in a trunk-and-drop system.

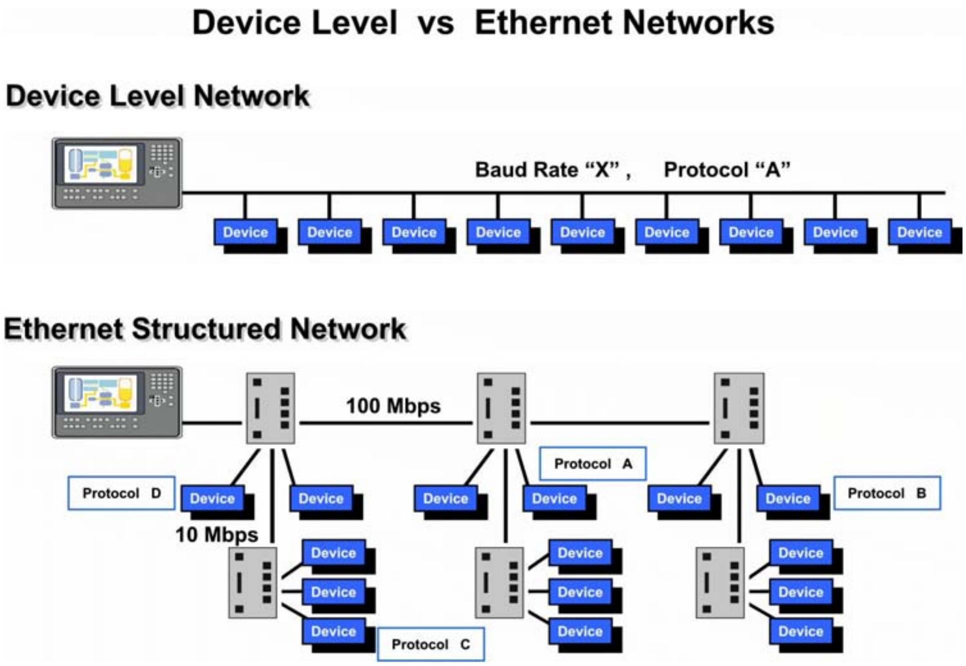
Device Level vs Ethernet Networks
Device Level Network
Baud Rate "X" , Protocol "A"
Device
Ethernet Structured Network
100 Mbps
10 Mbps
Protocol A
Protocol B
Protocol C
Protocol D
Device

Figure 3-1 Device-level Networks are Trunked. Ethernet Networks are Segmented.

These networks can connect more than 60 devices with communication cable without using any other components. Bus extenders or bridges increase the network length or add layout flexibility, but they are optional.

Industrial Ethernet uses “star” topology to connect infrastructure components to each other and to end devices. While industrial Ethernet requires that the cost and use of infrastructure components be considered in the total automation system design, the star topology provides almost unlimited flexibility for connecting to devices wherever they are located.

3.1.1 Implications of Multiple and Higher Data Rates

Device-level networks offer choices of several data rates versus distance options. Once these parameters are chosen for a particular network, all attached devices must be configured to operate at the same data rate. Industrial Ethernet, however, allows a mix of data rates within the same network, for example, 10 Mbps, 100 Mbps, and 1000 Mbps (1 Gbps). Higher speeds are generally used for connections (uplinks) between infrastructure components, such as switches and routers. For backward compatibility purposes, many industrial Ethernet devices are offered with both 10-Mbps and 100-Mbps capabilities, or all three speeds.

The ability to mix 10-Mbps, 100-Mbps, and 1-Gbps data rates in the same system allows engineers to balance the 1-Gbps benefits of maximizing throughput with the 10-Mbps benefits of higher immunity to electrical noise and lower sensitivity to custom cable and connector termination variances. The lower data rates of device-level networks allow long copper wire cable lengths of 500 meters (about 550 yards). With the higher data rates of industrial Ethernet, fiber-optic cable and/or infrastructure components are necessary for long distances.

3.1.2 Frame Size

A single frame of data on a device-level network is usually small (e.g., a CAN-based network contains only up to 8 bytes of data). A single frame of industrial Ethernet can contain up to 1500 bytes of data. The combination of high speed and high data capacity makes industrial Ethernet increasingly attractive as more intelligence is embedded into smaller and less expensive devices.

3.1.3 Implications of Multiple Protocols

Device-level networks are designed to communicate using one protocol. Even in cases where the same underlying technology is used for multiple networks (e.g., CAN technology for both DeviceNet and CANopen), placing one device of a different protocol on the network causes errors that prevent communication. However, IEEE Std 802.3-based Ethernet technology and infrastructure components handle multiple protocols.[6]

One advantage to open-standards-based Ethernet technology is that the user can choose a cabling approach and cable layout design, and select protocols no matter which vendors supply the products. Another advantage is the seamless connection between the plant floor, office, and Internet/Intranet worlds. With this connection, however, comes a challenge: keeping the enterprise and Internet network message traffic off the plant-floor control system. To prevent this traffic from slowing the industrial

[6] IEEE 802.3, IEEE Standard for Information Technology –Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Networks -- Specific Requirements – Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications *(Incorporates IEEE Std 1802.3D).* http://standards.ieee.org/getieee802/portfolio.html

network's response time, infrastructure components and addressing techniques provide connectivity while segregating plant-floor traffic.

3.2 Ethernet System Components

Industrial Ethernet systems require various infrastructure components to connect all the segments together. These include hubs, switches (especially managed switches), media converters, cables, and routers.

Classification of Ethernet Infrastructure Components

Ethernet infrastructure components are also known as Data Communications Equipment (DCE). Control devices that send and receive messages, such as variable frequency drives, PLCs, PCs, HMIs, machine vision devices, etc. are typically classified as Data Terminal Equipment (DTE).

3.2.1 Ethernet Hubs

Ethernet hubs, also referred to as repeaters or repeater hubs, are inexpensive electronic "repeaters" of Ethernet data packets; they have no built-in intelligence. They amplify the signal and can extend the distance and quantity of devices on a segment. All connected network devices hear all Ethernet data packets on all of the Ethernet hub ports. Also, all network devices connected to a hub must communicate at the same data rate.

Unlike switches, hubs only operate in a half-duplex communications mode, so they retransmit collisions to other devices. (See *4.3.1 Full-Duplex vs. Half-Duplex Implications*.) The cost of hubs initially was significantly lower than that of switches, so hubs were used as a lower-cost method of adding devices to the segments connected by switches. Switch products—thanks to home and office needs—have become ubiquitous, with prices dropping accordingly. Because switches do more and are relatively inexpensive, hubs have lost favor and are disappearing from vendor catalogs and store shelves.

Industrial switches, borrowing on the technology of their commercial counterparts, provide an attractive connect cost per device. It is possible to use switches (with all their performance benefits) in places where cost constraints formerly resulted in the use of hubs. It is recommended that switches be used throughout, and where economically feasible, managed switches (configurable) or routers should be used.

3.2.2 Ethernet Switches

Ethernet switches are intelligent connection devices that generally use high-speed hardware to buffer messages, analyze data frames, and connect the sending and receiving nodes in a virtual one-to-one connection. Switches are a fundamental part of most networks. They make it possible for several users to send information over a network at the same time without slowing each other down.

Ethernet Switch Terminology

Ethernet switches were historically referred to as bridges, multi-port bridges, or switching hubs. The simple term switch is now the industry's generally accepted term.

Switches are typically OSI Layer 2-type devices that allow all protocols to pass through. The Ethernet switch is the most effective means of setting up Ethernet systems, especially on machine-control systems on the manufacturing floor. In fact, the use of Ethernet switches is the preferred means of Ethernet interconnect of EtherNet/IP systems. Typical industrial switches can connect four or more network segments or Ethernet devices. Each of these connections (ports) typically supports both 10-Mbps and 100-Mbps data rates. Many of the newer switches will support 1-Gbps data rates as well.

3.2.3 Switches: Key System Elements

Switches are the key components that provide the determinism and throughput required for control applications. They do this through several key features:

- Switches learn which device addresses are connected to each of their several ports. They then send data received at one port only to the port that is connected to the target station. This reduces unnecessary traffic.

- If a switch is connected to a segment containing half-duplex devices (which may be connected via hubs), collisions from that segment are isolated to that segment and are not passed on to other devices or segments connected to other ports.

- The Ethernet connections (ports) on switches typically support and automatically configure themselves for either half-duplex or full-duplex communications. Using full-duplex communications from a switched port directly to a device turns off the traditional Carrier Sense Multiple Access with Collision Detection (CSMA/CD) mechanism so that collisions are impossible.

The switch’s ability to eliminate collisions is the most important mechanism to provide real-time capability for Ethernet-based control systems. If the data load increases on any one segment such that performance is affected, a switch may be added to split the data load between two or more segments, resulting in higher throughput. In addition, managed switches can prioritize traffic with up to eight traffic classes, allowing the preferential handling of real-time traffic over supervisory traffic (e.g., file transfer, web sessions, etc.).

3.2.4 How Switches Add Flexibility

The capabilities of switches increase the flexibility of the overall system in a number of ways:

- Switches can simultaneously communicate with and pass network messages between devices or segments operating at different data rates (10 Mbps, 100 Mbps, or 1 Gbps).

- Provided with both types of ports, they can connect to both fiber-optic and twisted-pair cables, allowing developers to use a mix of media in the same system.

- Managed switches come with added functions. Each has its own IP address, through which users can enable/disable or configure the added functions. The features included with a managed switch vary from vendor to vendor. Some vendors offer managed switches that provide added diagnostics, addressing functions, and redundancy. Others can be configured to help manage traffic through configuration optimization of Ethernet features (e.g., virtual LANs, Quality of Service), bandwidth allocation, user authentication, and other sets of features.

Managed Switches

Managed switches provide the basic switching capabilities of unmanaged switches, plus they offer additional diagnostic or traffic-management functions.

In essence, a managed switch is configured or tailored for a particular part of the overall application. The different switches in an application can each have different configurations that the user must maintain. If a managed switch is replaced, the particular configuration settings must be downloaded or entered into the new switch. Unmanaged switches, however, perform only the basic switch functions and do not need to be reconfigured if they are replaced. Since managed switches perform the basic switching functions "out of the box," they can be used as "plug-and-go" replacements for unmanaged switches. For more specific details on the requirements for switches in EtherNet/IP networks, see *Appendix A Recommendations for Ethernet Switches in EtherNet/IP Systems*.

3.2.5 Media Converters

Hubs and switches that allow connection of fiber and twisted-pair wire to the same device may be purchased. In cases where the mix of ports does not match the requirements, twisted-pair to fiber-media converters are available for converting a single cable of one type to another.

3.2.6 Routers

Routers provide added functionality not usually found in switches. A hub or a switch will pass along any broadcast or multicast packets they receive to all the other segments, but a router will not. A four-way traffic intersection is a good example. With a hub or unmanaged switch, all of the traffic passes through the intersection no matter where it is going. If, however, this intersection is at an international border, the rules are different. To pass through the intersection, the driver must provide the border guard with a specific destination address. If the driver doesn't have a specific destination address, the guard will not let the vehicle pass.

Routers work the same way. Without a specific address of a destination device, the router will not let the data packet through. This is good for keeping networks separate from each other, but is less desirable when different parts of the same network need to communicate. Switches allow different devices on the same network to talk to each other.

Routers are typically used when traffic from two different networks needs to be interconnected. Routers direct traffic based on the network address (OSI Layer 3) information of the packet. Routers can help the performance of factory-floor Ethernet networks by blocking Ethernet multicast and broadcast data traffic. Routers can isolate the traffic of other connected but unrelated Ethernet networks, which could otherwise reduce the performance of factory-floor systems. Routers also can determine the optimal path along which network traffic should be forwarded and can be used to help implement network security.

Routers normally do not have as many interfaces (ports) as switches, but like managed switches, they require setup and configuration. Most switches add very little latency (or delay) to the messages that pass through them. However, based on the added functions a router performs and the specific vendor's hardware implementation, added latency may be introduced. Engineers need to check the latency specifications of a router if real-time control information will pass through it.

3.3 Ethernet Cabling Systems

New Ethernet installations are based on twisted-pair and fiber-optic cables. The following presents general guidelines and an overview of cabling schemes. The use of industrial grade Ethernet cable is critical for withstanding exposure to harsh environments.

3.3.1 Star Topology: The Standard Approach

Industrial Ethernet systems are laid out in a point-to-point cabling scheme with one cable used for each device being connected. These individual connections are coordinated through an Ethernet switch in a cabling topology known as a star configuration.

The star configuration simplifies troubleshooting for device, cable, and connection installation problems. With this star approach, multiple infrastructure components are interconnected, and each infrastructure component connects in a star fashion to a group of individual control devices. This approach provides a flexible, distributed control architecture.

Most manufacturers offer pre-made or custom “patch” cables in a wide variety of lengths. Typically a backbone of switches will be connected with copper cables up to 100 meters (328 ft) in length or longer fiber-optic cables, where each switch isolates a machine or major part of a machine. The other ports of the switch connect to the control devices using twisted-pair or fiber cables for that part of the machine or manufacturing process.

In a situation where there are more devices than switch ports to accommodate them, connecting one of the first switch’s ports over to a port on a second switch effectively adds more ports to the first switch. Some active infrastructure components have a modular hardware-expansion capability where additional ports may be added (without using a device connection port) by plugging together port expansion modules. Another option is to select a switch with additional ports and use the smaller switch elsewhere. It is often wise to choose a larger switch than necessary and leave a few open ports. This way, when it comes time to add sensors, actuators or subnetworks, it will not be necessary to replace the switch.

3.3.2 Implementation of Mixed Cable Systems

Infrastructure components, such as media converters, allow copper wire (twisted-pair) and fiber-optic cables to be mixed in the same system. Industrial Ethernet systems may require the use of fiber optics to cover the longer cable-run distances that would previously be implemented with copper wire at lower data-rate settings in device-level networks. Also, the added noise immunity of fiber, especially at the 100-Mbps and 1-Gbps rates, can reduce start-up surprises. This is especially true when routing wires near welders, larger horsepower drives, motors, or motor starters.

3.3.3 Applicable IEEE Standards for Twisted-pair Ethernet

The major IEEE standards that apply to industrial twisted-pair Ethernet installations are 10 Base-T for 10-Mbps connections and 100 Base-TX for 100-Mbps or fast Ethernet operation. The term “fast Ethernet” applies to the standards that support 100-Mbps operation. Gigabit Ethernet (1000 Base-T) operates at 1000 Mbps.[7]

3.3.3.1 Unshielded Category 5e Cable

The Unshielded Twisted Pair (UTP) wiring standard, which allows for flexibility and ease of installation and maintenance, is the most prevalent cabling scheme in use today. There are two twisted-pair (four

[7] IEEE Std 802.3. http://standards.ieee.org/getieee802/portfolio.html

wires) and four twisted-pair (eight wires) UTP cables. Two twisted-pair cables support network speeds of 10 Mbps as well as the 100-Mbps. Industrial Ethernet systems with 1-Gbps speeds use four twisted-pair cables with a maximum cable length of 100 meters (328 ft). Pre-configured Category 5e (CAT 5e) cables are available in required lengths. Field terminations may require manual work with crimping tools, bulk cable, and connectors.

The quality of the cable and of the connector termination is particularly important at the high data rates and in the electrically noisy environments of industrial Ethernet.

Only the highest quality, industrial grade Category 5e (or higher) cable should be used in industrial Ethernet applications. Category 5e (*e* for *enhanced*) has superseded the older CAT 5 designation for 100-Mbps systems. CAT 5e can contain either two twisted pairs (four wires) or four twisted pairs (eight wires). The added wires and improved signal transmission characteristics allow CAT 5e cables to be used for 1-Gigabit Ethernet and all power-over-Ethernet applications. However, two twisted-pair cables have the advantages of compatibility with M12 (IP65 and IP67 grade) connectors and, in 100-Mbps applications where only two pairs carry communications, this reduces the chance of electrical noise interference.[8]

3.3.3.2 Shielded Category 5e Cable

Shielded cable should be considered for use in high noise environments. It also is the prevalent cable used in European installations. See *4.4 Layout and Wiring Considerations* for more information.

3.3.3.3 Connectors and Pre-terminated Cables

The RJ-45 connector is the standard for twisted-pair cabling. The IEEE Std 802.3[9] and TIA Std TR42.9[10] industry standards groups are in the process of specifying IP65 and IP67 versions of industrialized Ethernet connectors.

“D”-coded M12 connectors with four wires are used to create IP65 and IP67 fluid-resistant, sealed connections. These connectors allow removable connections to panels and connection of machine-mount control devices. See *7.1.2 M12 “D”-Coded Connector for EtherNet/IP* for more information.

The use of pre-terminated cables should get serious consideration, especially for 100-Mbps and 1-Gbps connections. Improperly terminated cables can degrade the signal or provide a path for electrical noise to enter, causing intermittent communications or reduced system throughput as the devices are forced to re-transmit a higher percentage of messages.

3.3.3.4 Crossover vs. Straight-through Cables

Unlike device-level networks, different types of Ethernet devices require different pin-to-pin connections. Some devices require the use of crossover cables where the transmit pins at one end of the cable are connected to the receive pins of the other connector. Other devices require a 1:1 (straight-through) cable where the transmit pins at one end are connected to the transmit pins of the other.

[8] The Ingress Protection (IP) code IP65 refers to equipment that is dust-tight and protected against water jets. IP67 refers to equipment that is dust-tight and protected against immersion in water. See *Appendix E* for more information.

[9] IEEE 802.3, http://standards.ieee.org/getieee802/portfolio.html

[10] TIA TR42.9, Industrial Telecommunications Infrastructure Standard. www.tiaonline.org

Recommended Cable Practice

Different cable types can be distinguished through the use of different colored outer cable jackets. Most infrastructure components automatically compensate for the type of device connected. This is known as auto-crossover or auto MDI/X capability.

There are two classes of Ethernet devices: Data Communications Equipment (DCE) refers to infrastructure components, and Data Terminal Equipment (DTE) refers to control devices. The following connection schemes normally apply:

- **DCE to DCE** (infrastructure components connected to infrastructure components): crossover cable
- **DTE to DTE** (control devices connected to control devices): crossover cable
- **DTE to DCE** (control devices connected to infrastructure devices): 1:1 straight-through cable

3.3.3.5 Other Options

Like device-level networks, a variety of cable options are available for easier wiring in specific application situations (e.g., in cabinets versus outside of cabinets). See *4.4 Layout and Wiring Considerations* for more details.

3.3.3.6 Future Options

The IEEE Std 802.3af task group is working on DTE powering, which will enable switches and other DCE equipment with integrated power supplies to provide dc power for unpowered DTE equipment.[11] Power will be carried either by signal pairs 1 & 2 and 3 & 6, or by unused pairs 4 & 5 and 7 & 8. This capability can provide power to sensors and video equipment as well as security devices and wireless access points, extending the flexibility of subnetworks.

3.3.4 Applicable IEEE Standards for Fiber-Optic Ethernet

Fiber-optic cable mediums provide longer distance runs than wire as well as immunity from electrical and radio frequency (RF) noise. The major IEEE standards that apply to industrial fiber-optic Ethernet installations are 10 Base-FL for 10 Mbps connections and 100 Base-FX for 100-Mbps (fast Ethernet) operation; 1000 Base-SX refers to IEEE Std 802.3z for 1-Gbps (or Gigabit Ethernet).[12]

3.3.4.1 Glass Fiber Cable

The fiber-optic cable typically used in both 10-Mbps and 100-Mbps industrial Ethernet applications is based on 62.5-micron duplex multimode glass cable. Also available are 50- and 100-micron fibers. Maximum cable lengths are based on the type of cables used and whether full-duplex or half-duplex

[11] IEEE 802.3, http://standards.ieee.org/getieee802/portfolio.html

[12] IEEE 802.3, http://standards.ieee.org/getieee802/portfolio.html

communications is used. Cable lengths of 2000 meters (6562 ft) are possible using multimode fiber in full-duplex, 100-Mbps applications. The use of pre-terminated glass fiber-optic cables is recommended.

Fiber-optic cables are terminated using SC, ST, or MTRJ connectors. Like twisted-pair wiring, the use of half-duplex communications has additional restrictions as to how many hubs can be used between the sending and receiving devices. These issues are covered in more depth in *4 Designing the Infrastructure*.

3.3.4.2 Plastic Fiber Cable

For cabling within cabinets or over short distances in conjunction with media converters, plastic fiber optics provide noise immunity and easier connection capability, especially if custom cables need to be used. Polymer (980/1000 um) fiber cables can be used for lengths up to 50 meters (164 ft), while HCS (200/230 um) fiber cables can be used for up to 100 meters (328 ft).

3.3.4.3 More Cabling Resources

An overview of characteristics to consider when selecting connectors and cables suitable for industrial applications can be found in *7 Selecting Components*. For complete connector, cabling, and layout information, refer to the *EtherNet/IP™ Media Planning and Installation Guide* available from ODVA.[13]

3.3.5 Wireless Ethernet

Wireless Ethernet network systems use radio frequency (RF) transmitters and receivers to send and receive data. Like any system based on radio, performance and interference issues need to be addressed for use on the plant floor. Wireless systems are based on the IEEE Std 802.11b and other standards; however, this remains an emerging area for industrial Ethernet applications and is not covered here.

3.4 Ethernet Data Connections

Ethernet data packets are sent in the format shown in Figure 3-2.

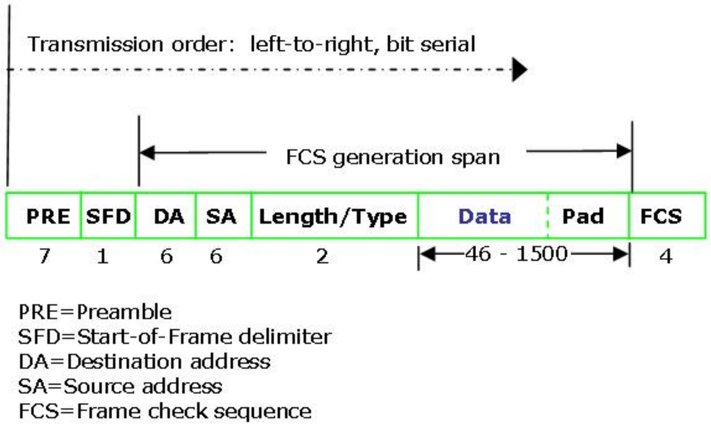

Figure 3-2 Typical Layout of an Ethernet Data Packet.

This data format is used to implement the Media Access Control (MAC) protocol that allows a device to "talk" on the Ethernet network. Each MAC device has a unique Source Address (SA) comprised of a 6-byte number (48 bits or 12 hexadecimal digits) that was assigned to it at the time of manufacture. The

[13] ODVA, Inc., *EtherNet/IP™ Media Planning and Installation Guide*, Ann Arbor: 2006. www.odva.org

Destination Address (DA) is the target MAC address for which the packet of data is intended. Setting the first bit to a "1" in the DA field, indicates a packet of data for multiple destinations. This enables an Ethernet device to transmit one packet that can be received by multiple other devices.

There are a number of different types of Ethernet packets that can be sent and received on an Ethernet network. Some of these protocols are Novell's IPX/SPX, DECNET, UDP, TCP/IP, FTP, TELNET, and so on. All of these unique protocols use the MAC to do the physical sending and receiving of data packets. However, by defining how the "DATA" portion of the data packet is organized, different protocols and functions are created.

3.4.1 IP Addressing of Industrial Ethernet Devices

One of the more common protocols familiar to many is TCP/IP, which stands for Transmission Control Protocol/Internet Protocol. EtherNet/IP uses the TCP/IP suite to send messages between devices. The TCP portion of this protocol suite allows for reliable connections on the Ethernet and the quick streaming of data from one Ethernet node to another. It has become the most popular protocol implemented on Ethernet networks.

The Internet Protocol (IP) portion of the TCP/IP suite is its routing mechanism. All devices on an industrial Ethernet-based system must be assigned an IP address. Infrastructure components with added diagnostic capabilities (managed switches and hubs) also require an assigned IP address. It is most commonly identified by the 4-byte address listed in the network properties on personal computers that use TCP/IP as their Ethernet network connection. (See *2 Understanding the Basics of Network Protocols* for examples). IP addresses create unique, high-level identifiers for IP networks, such as the Internet.

Media Access Control (MAC) Address

Unlike device-level networks where the physical layer address also is the address that the user sets, Ethernet devices are addressed by assigning an IP address to that particular device's MAC address. The MAC address is the lowest level address and resides at the data link layer for Ethernet.

Each Ethernet device vendor sets a unique MAC address into each product that is shipped. Before starting up the system, the engineer must first assign an IP address to a plant-floor Ethernet device as a part of its configuration. See *4.6 Performance Considerations: The Need for Switches* for more information.

3.4.2 Important Ethernet Tools and Protocols

EtherNet/IP uses the TCP/IP suite as the basis for sending and receiving the Common Industrial Protocol (CIP) packets of data. CIP is the protocol used by DeviceNet, ControlNet, and EtherNet/IP to facilitate the I/O control, programming, startup, operation, and maintenance of control systems. By encapsulating these CIP messages inside the TCP/IP Ethernet protocols, EtherNet/IP is able to provide a powerful control system network that can be seamlessly integrated into the business systems of a manufacturing facility. Control engineers should become familiar with a number of other protocols associated with the TCP/IP suite:

- UDP - User Datagram Protocol
  UDP packets offer a simple method of moving data with minimal overhead. UDP is used by EtherNet/IP to move connected I/O messages in a very efficient producer-consumer model.

- DHCP - Dynamic Host Configuration Protocol
  This is a protocol to allow a device, acting as a DHCP server, to assign 4-byte IP addresses to devices connected to the network.

- DNS - Domain Name Services
  This protocol searches for resources on a TCP/IP network using a database connected to the TCP/IP network. It makes correlations between computer names, their 4-byte IP addresses, and the resources and information available in the network. There are many other tools and protocols available. See *5 Deploying the Network* for more information.

3.5 Diagnostics via Infrastructure Components

Industrial Ethernet infrastructure components, like their device-level network counterparts, can provide a full range of diagnostics over displays such as built-in LEDs or seven-segment readouts. Typical diagnostics include transmit/receive status by port, port active/cable connected status, and 10-Mbps versus 100-Mbps operation. Most intelligent infrastructure components perform periodic self-tests and provide LED indication of failure. Some units also provide a hardware contact that can be wired to an alarm or nearby industrial controller to signal a device failure.

3.5.1 Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is an industry standard protocol supported by virtually all infrastructure component manufacturers. Through the SNMP protocol, engineers can set configuration parameters (e.g., assign IP addresses) and receive diagnostics (i.e., remote monitoring data). This information, however, only relates to infrastructure components, not to Ethernet embedded controllers, devices, or I/O. EtherNet/IP is used to access Ethernet control device data.

SNMP drivers are available as standard in virtually all infrastructure-related (network management) software packages. SNMP allows a software package from vendor X to display and set parameters in vendor A, B, and C’s infrastructure components. Managed infrastructure components provide added diagnostics, such as port status (cable connected/unconnected, communications active/inactive), packet size, and statistics on the packet type transmitted (unicast, multicast, broadcast), power supply status, etc. These devices contain a data table called a Management Information Base (MIB) that contains diagnostic and other information.

The SNMP defines the format and standard types of remote monitoring (RMON) information. This standard allows multiple vendor diagnostics to be accessed by vendor-independent software. The standard also allows vendor-specific information to be added. Within a point-to-point tree structure containing hundreds of connected devices, the ability to quickly locate cabling or other faults in the Ethernet infrastructure can be key to expediting start-ups. For more information, see *5 Deploying the Network*.

3.5.2 Traps

Managed infrastructure devices allow a network engineer to set up SNMP traps. If it is important that multiple devices receive diagnostic information, all the extra network traffic could significantly affect the response time of the control network. Traps concentrate, or “trap,” the diagnostics from a quantity of infrastructure devices. Without traps, PCs, etc. would have to constantly poll the MIB data from all the infrastructure devices to determine if errors occurred. Infrastructure components can be configured with the IP addresses of one or more networked PCs and devices.

When an alarm or error condition happens (some infrastructure diagnostics can determine if a parameter exceeds a high or low limit), the infrastructure component sends an error message to the trap. This concept is similar to the “Change of State” messaging found in CIP. An event, in this case an error, triggers a message to a specific device, in this case a trap. This “report by exception” approach limits Ethernet network management traffic, allowing greater network bandwidth to be used for control.

3.5.3 Network Management and Web-based Management Software

Many vendors offer network management software that uses SNMP, Bootstrap Protocol (BootP), and accessed MIB data to allow an operator to view the status of the entire network. It also allows the IP addresses to be loaded into the devices. This overview of the installation can be used to identify problem areas during startup or on an ongoing basis. Detailed information about the devices may also be accessed.

Many devices have an integrated Web-server that allows SNMP data to be accessed using standard Internet browsers. The combination allows problem areas to be identified using network management software with the Internet browser to “zoom in” on the device’s web page(s) for access to detailed information or settings. Some devices also include embedded data sheets or user manuals that can be accessed.

3.6 Isolated vs. Integrated Networks

Historically, device-level networks have been “isolated” from one another. A single controller, such as a PLC or PC, is the “master” of a small network of nodes, such as I/O blocks, variable frequency drives (VFDs), barcode readers, servo drives, and so on. This small control network does not have a physical connection to any other network other than through the master PLC or PC. As such, communications between these device-level networks require a separate network port on the PLC or PC and also require programming in the controller to send and receive data. In this way, the network communications from machine to machine, from cell to cell, and throughout a manufacturing campus have been isolated to only what was needed for the specific connected devices.

These isolated networks, while very effective at providing distributed control, do not provide the “seamless” connectivity that industrial Ethernet offers. The ability to connect and communicate between the business system networks and the plant-floor networks via industrial Ethernet provides a huge opportunity to reduce work-in-process inventory, improve throughput, find manufacturing bottlenecks, increase uptime, and improve customer service. The “integrated network,” while providing these manufacturing improvement opportunities, demands a conscientious effort at an overall network design and layout for an entire facility. The benefit of Ethernet is that it can be scaled, from small original equipment manufacturer (OEM) machine control applications, to integrated manufacturing lines. As the applications become more complex, the quantity, type, and functionality of the required infrastructure components also grows. (See *6 Infrastructure Application Scenarios* for more information.)

Typically, manufacturing and OEM control engineers can achieve the goals of a good industrial Ethernet network design by using robust industrial Ethernet components. First, the industrial Ethernet electronics and cabling systems should implement full-duplex switching technology, follow proper Category 5e (or higher) twisted-pair and fiber installation techniques, and address the environmental conditions present around the equipment. Next, using the proper equipment and building industrial Ethernet control networks in a “managed network” scheme, manufacturing and OEM control engineers can achieve machine control objectives while adding the benefits of total plant network integration. Finally, when integrated networks with industrial Ethernet are used, new requirements such as security and network traffic routing can arise. Additional “higher-level” network equipment and networking techniques may be required to meet these needs.

3.6.1 Security Issues

Firewalls are important to the security of integrated networks. Routers and switches may also be configured to contribute to network security.

3.6.1.1 Firewalls

A firewall is a software and hardware solution that acts as a barrier to protect a network from outside intruders when connected to a wide area network (WAN), such as an enterprise network or the Internet. Any plant-floor information (e.g., remote diagnostics) that is accessed from outside the building must first pass through a firewall.

3.6.1.2 Routers and Switches

Most routers today allow filtering of TCP/IP packets as they approach the internal network from a WAN or the Internet. Operating at the network layer of the OSI Model, these routers can deny access to the internal network for packets from specific external sites. They allow access to only certain services (such as internal web servers) or certain computers in the internal network (such as a mail server). All attempts to connect to other computers in the internal network will be denied. Although not bulletproof, routers do provide an added line of defense. Many managed switches also contain the ability to disable unused physical ports to prevent unauthorized personnel from plugging in and accessing the network.

3.7 Traffic Management with VLANs

A virtual LAN (VLAN) is a switched network segmented by functions or applications on an organizational basis as opposed to a physical or geographical basis. VLAN addressing mechanisms allow stations to be assigned to logical groups that communicate across multiple LANs as though they were on a single LAN.

Bridges and switches filter Destination Addresses and forward VLAN frames only to ports that serve the VLAN to which the traffic belongs. One example is a device-level network with “multi-master” capability. In this case, two industrial PCs communicate on the same network with a number of devices. PC #1 may control devices 1, 3 and 5 while PC #2 controls devices 2, 4 and 6.

With industrial Ethernet, all eight devices can be connected together on the same Ethernet network using several switches. A VLAN approach might place PC #1 and devices 1, 3 and 5 on VLAN 1, and PC #2 with devices 2, 4 and 6 on VLAN 2. Devices or network segments related to VLAN 1 do not receive messages from VLAN 2, thus reducing overall network traffic.

3.7.1 How VLANs Work

A VLAN is created by “tagging” or inserting a 4-byte VLAN header into the basic Ethernet (MAC Data) frame between the Source Address and Length/Type fields as shown in Figure 3-3.

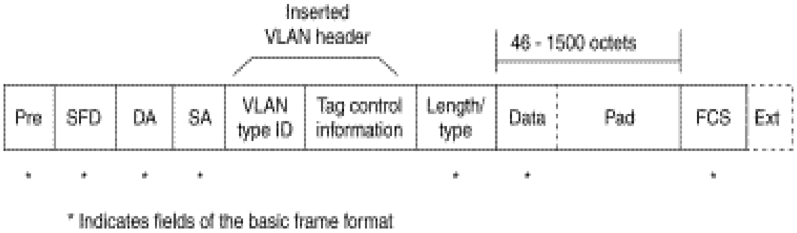

Figure 3-3 VLAN-tagged Frames are Identified when the MAC Finds the LAN Type Value in the Normal Length/Type Field Location.

The VLAN header contains a Tag Control field that defines the transmission priority (0 to 7, where 7 is the highest) and a VLAN ID that identifies the particular VLAN over which the frame is to be sent. These provide the means to expedite time-critical network traffic. Ethernet switches or routers use the VLAN header information to direct the message to the correct port and transmit the frames according to their priority. VLAN tagging requires that all infrastructure devices involved with a VLAN group be equipped with the VLAN option.

3.7.2 VLAN Benefits

The application of VLANs allows the flexibility of using interconnecting devices while reducing unnecessary network traffic and delay time. In addition to managing network traffic, VLANs provide added system scalability and security, and they simplify network management by making device adds, moves, or changes easier to administer. The configuration of a VLAN requires managed infrastructure components.

3.8 Topologies to Provide Ethernet Network Redundancy

Infrastructure components come with a variety of system redundancy options, from redundant power supplies, to cable and communication path redundancy. Cable and communication path redundancy is often required in larger scale installations or in cases where an infrastructure component or cable failure may cause significant scrap of manufactured product or result in damage to a machine or process.

There are both standardized and proprietary methods of achieving redundancy in Ethernet networks that involve variations on standard Ethernet topologies. Standardized methods have the benefits of allowing the mix of multi-vendor products or the ability to easily change vendors without re-engineering the system. Like PLCs, proprietary redundancy methods may offer special features or performance advantages in certain applications, but the approach cannot be mixed with multiple vendors' components. In addition, special attention needs to be paid to how, and if, the proprietary method can be used in conjunction with standardized redundancy approaches.

If a user mixes proprietary and nonproprietary methods within an overall system, communication loops can be formed, creating spurious traffic that can overload and effectively shutdown the network. Different vendors' equipment should be evaluated based on the type of redundancy offered, the switchover time, and if there is message storage capacity to reduce "lost messages" during switchover time.

3.8.1 IEEE Standard Network Redundancy Options

The Spanning Tree and Trunking/Aggregation methods are standard, multi-vendor approaches to achieving network redundancy.

3.8.1.1 Spanning Tree Protocol Method

The Spanning Tree Protocol (STP) is a Layer 2 (data link layer) protocol designed to run on bridges and switches. The specification for STP is IEEE Std 802.1D.[14] The main purpose of STP is to ensure that there are redundant paths while avoiding a loop situation that causes network messages to endlessly cycle around the loop and repeat, causing the network to saturate and communications to effectively stop.

As shown in Figure 3-4, switches with STP can detect one or more cable cuts and automatically determine the next best path for the data to travel. This capability requires switches that support the Spanning Tree algorithm, and transmit and receive Spanning Tree configuration messages. Vendors should be consulted to determine how many bridges/switches can be supported, and to determine the length of the time lapse from when a break is detected to when the network is reconfigured and normal communications resume.

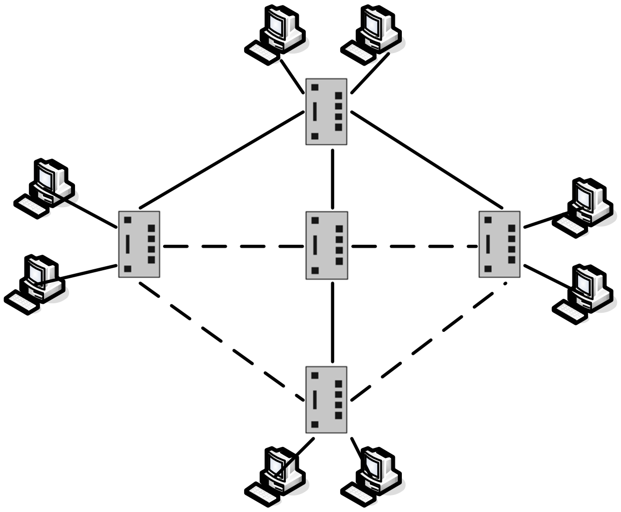

Figure 3-4 Using the Spanning Tree Protocol Method, Switches Detect a Cut Cable and Automatically Determine the Next Best Path for the Data to Travel.

[14] IEEE 802.1D, IEEE Standard for Local and Metropolitan Area Networks – LAN/MAN Bridging and Management – Media Access Control (MAC) Bridges. http://standards.ieee.org/getieee802/802.1.html

The Rapid Spanning Tree (RSTP) specification (originally IEEE Std 802.1W, now included in IEEE Std 802.1D) is an enhancement over the original Spanning Tree specification. It allows for faster recovery times (a few seconds) than Spanning Tree. In Spanning Tree applications, recovery times typically range from 30 to 60 seconds. Large and complex topologies may result in minutes of recovery time.

3.8.1.2 Link Aggregation

The Link Aggregation method (IEEE Std 802.3AD[15]) is used to provide redundant communication paths between two devices when two or more links connect the devices as shown in Figure 3-5. When one link is broken, the others carry the traffic.

Link Aggregation uses all the links all the time, typically in situations where high traffic volume (such as uplinks from lower-level networks to higher-level enterprise networks) requires the added bandwidth of multiple communication paths. For example, four fast Ethernet links (100 Mbps) produce a bandwidth of a single 400-Mbps link. In a case where one of the links is cut, traffic continues but at a slower 300-Mbps equivalent bandwidth.

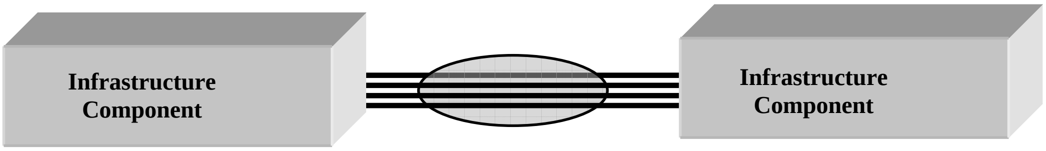

Figure 3-5 When One of Four Fast Ethernet Links between the Two Devices is Broken, Message Traffic Capacity Decreases but Continues.

3.8.2 Ring Redundancy

Ring redundancy schemes are proprietary in nature and are available from multiple vendors. These redundancy schemes are not interoperable between different vendors' infrastructure products. Redundancy configurations require multiple switches to be connected in a ring manner as shown in Figure 3-6. These systems allow typical communications recovery times in the 100s of milliseconds range.

[15] IEEE 802.1AD, IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks. http://standards.ieee.org/getieee802/802.1.html

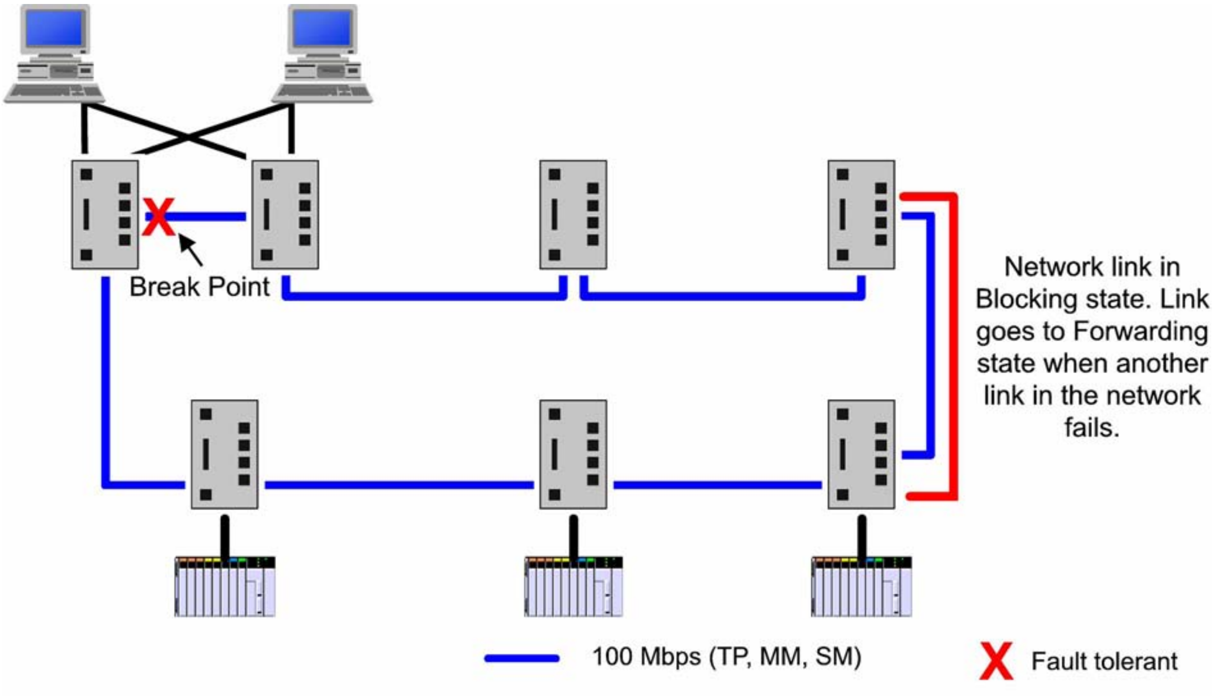
Break Point
Network link in Blocking state. Link goes to Forwarding state when another link in the network fails.
100 Mbps (TP, MM, SM)
X Fault tolerant

Figure 3-6 Ring Redundancy Topology.

3.8.3 Dual Homing

The dual homing approach is typically used to achieve redundancy between routers or network layer (Layer 3) switches that have router-level functions. (See *4.11 Multilayer Ethernet Switches* for more information.) In addition to a simple hub-and-spoke network where a remote router is connected to a single distribution router, the remote router can be dual-homed to two or more distribution routers. A dual-homed remote router will have two or more distribution routers.

Figure 3-7 shows a common dual-homed remote topology with one remote router; however, 100 or more routers could be connected to Distribution Routers 1 and 2. The remote router will use the best route to reach its destination. This approach typically allows flexible configurations, combinations of two ports with different data rates (10 Mbps, 100 Mbps, or 1000 Mbps) and different media (fiber optic, twisted pair), and typically faster switchover times than the spanning tree method.

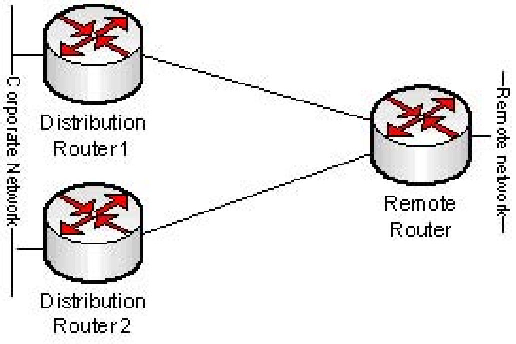

Figure 3-7 If Distribution Router 1 Experiences a Failure, the Remote Router Can Use Distribution Router 2 to Reach the Corporate Network.

4 Designing the Infrastructure

Network design is a methodical, step-by-step process. It includes the top-down process of mapping out the network infrastructure required for an application. This systematic process focuses on the application, technical objectives, and business goals.

When designing the infrastructure for an EtherNet/IP network, it helps to first define a logical view of the network, including a description of traffic-flow and architectural topology, before designing the physical layout. With a top-down network design, the emphasis is put on planning before implementation. If the system involves a fully integrated control and business network, cooperation between control and Information Technology personnel cannot be overemphasized. The two groups need to work together. What follows is a review of the processes and steps involved in network design. This approach considers both logical and physical network design, takes an in-depth look at infrastructure components, and provides some guidelines for practical network design.

4.1 Network Design Method

The network design process can be viewed as a sequence of the following steps:

1. Definition of application requirements
2. Logical network design
3. Physical network design

The following sections list activities recommended for consideration during EtherNet/IP network design, including functional application requirements, project scope, and other factors. To better understand the design issues, they are presented in the form of a comparison between industrial networks with active infrastructures (such as EtherNet/IP) and passive infrastructures.

4.1.1 Design Steps for Isolated EtherNet/IP Networks

Except for the active infrastructure of EtherNet/IP, the design process of an isolated EtherNet/IP network is similar to a device-level network as shown in Figure 4-1.

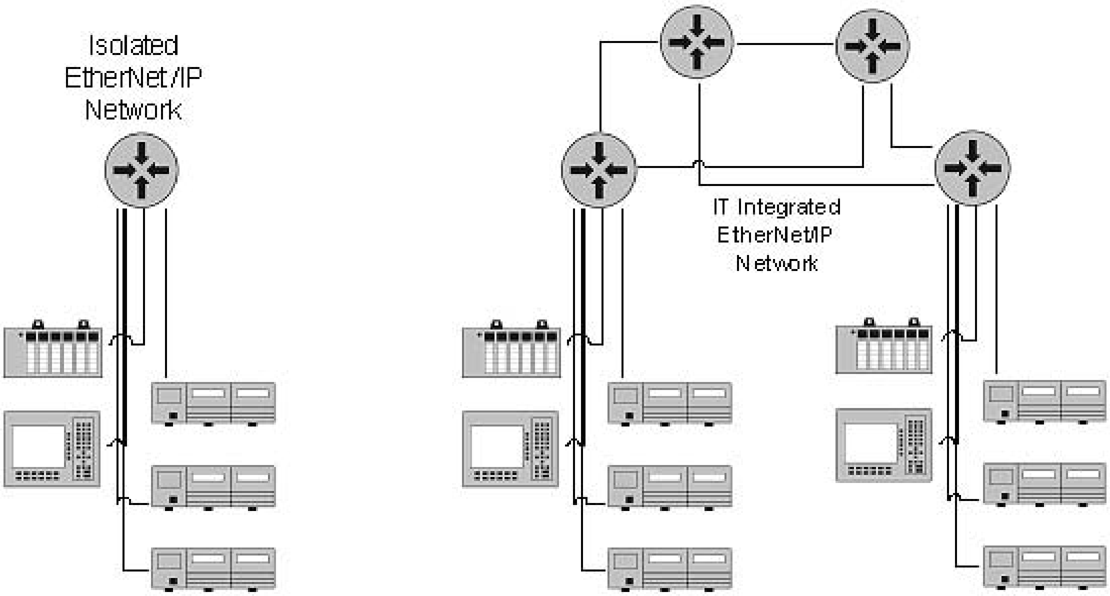
Isolated EtherNet/IP Network
IT Integrated EtherNet/IP Network

Figure 4-1 Comparison of Isolated and Integrated EtherNet/IP Networks. (Courtesy of General Motors.)

An isolated network has no immediate, continuous connection with the rest of the plant floor network or the enterprise network. Typically, it will connect to a supervisory PC or node, which will then act as a gateway to filter out connections from/to the main network. The design process for an isolated network is broken into three steps, which follow.

4.1.2 Definition of Application Requirements

During this stage, application requirements are identified and analyzed. The collected data is then used as design input. See Table 4-1. Note that EtherNet/IP networks will require more attention to planning and design than proprietary industrial device-level networks, which have many predefined specifications.

Table 4-1 Definition of Application Requirements

Activity	Industrial Network with Passive Infrastructure	EtherNet/IP
Define project goals and constraints (e.g., budget constraints).	This activity is usually performed. Controls company may assist.	This activity should be performed. Bring in IT Dept. and outside help for expertise if necessary.
Analyze technical goals and constraints (e.g., scalability, performance, and device configuration requirements).	This activity is usually performed. Controls company may assist; may be part of overall control scheme built around proprietary systems.	This activity should be performed. Bring in IT Dept. and outside help for expertise if necessary.

4.1.3 Logical Network Design

Logical network design is a process of creating a logical model of a network infrastructure consisting of the steps listed in Table 4-2. Begin by designing a network topology.

Table 4-2 Logical Network Design

Activity	Industrial Network with Passive Infrastructure	EtherNet/IP
Design network topology.	This activity is normally not performed because network topology is predefined.	This activity should be performed. Enlist help of IT Dept., system integrator, or control engineer familiar with Ethernet systems.
Define IP addressing and naming.	A node-addressing model is usually created. Node naming is usually not supported in these networks.	This activity should be performed. May need help from IT Dept.
Develop policies/procedures for network operation and maintenance.	Typically performed by plant engineers and maintenance personnel.	May be performed by either plant or IT/MIS personnel.
Define switch and router features.	Not applicable in the context used in information networks.	This activity should be performed. Get IT help if necessary.

4.1.4 Physical Network Design

Once application requirements have been determined and a logical network design has been achieved, it is time to design the physical network. See Table 4-3. This stage may be more complex in working with EtherNet/IP, but it will pay off in the long run if plant-floor connectivity is needed.

Table 4-3 Physical Network Design

Activity	Industrial Network with Passive Infrastructure	EtherNet/IP
Select cable system components and infrastructure devices.	This activity is equivalent to selection of the network cable system components.	This activity must be performed. System integrators and IT Dept. may help.

4.2 Non-Isolated EtherNet/IP Networks

A non-isolated network as shown in Figure 4-1 is integrated with the rest of the plant floor network and can communicate upwards to the enterprise network. Since there are more integration concerns (e.g., working with corporate IT), initial design work will be necessarily more involved. The design steps follow.

4.2.1 Definition of Application Requirements

During this stage, application requirements are identified and analyzed (Table 4-4). The collected data is then used as design input. Note that an extra characterization step has been added.

Table 4-4 Definition of Application Requirements

Activity	Industrial Network with Passive Infrastructure	EtherNet/IP
Define project goals and budget constraints.	This activity is usually performed but has not been recognized as a separate activity.	This activity should be performed, if appropriate. Work with IT Dept.
Analyze technical goals and constraints (e.g., scalability, availability, performance, security, manageability, usability, and adaptability).	This activity is usually performed. Work with controls company.	This activity should be performed. Get outside help if necessary. Work with IT Dept. or system integrator.
Characterize the existing network.	This activity is rarely performed since these networks are normally designed from scratch.	It is expected that these networks will normally be designed from scratch. Therefore, this activity will be rarely performed. An interface with the existing plant/enterprise network should be characterized instead.
Characterize network traffic. This activity may include characterization of traffic flow, traffic load, traffic behavior, and Quality of Service.	This activity is usually performed as part of network planning activities. Includes transaction analysis.	Both transaction and traffic analysis should be performed. Attention should be paid to multicast traffic generated by implicit messaging on EtherNet/IP.

4.2.2 Logical Network Design

Logical network design (Table 4-5) is a process of creating a logical model of a network infrastructure consisting of the steps listed in the table below.

Table 4-5 Logical Network Design

Activity	Industrial Network with Passive Infrastructure	EtherNet/IP
Design network topology.	This activity is normally not performed because network topology is predefined. For example, proprietary networks use a fixed and highly-defined topology.	This activity should be performed. It is recommended that network topology be consistent with plant/enterprise network topology. It is also recommended to address here, if applicable, how the EtherNet/IP network under design will fit into a typical three-layer hierarchical model of a building/campus network with need for path redundancy, load balancing, VLANs, and physical security.
Design a model for addressing and naming.	A node-addressing model is usually created. Node naming is usually not supported in these networks.	For a non-isolated network this activity is usually not performed because node naming and addressing must be the same as in the plant/enterprise network to which this network will be connected or integrated.
Define switch and router features.	Not applicable in the context used in information networks.	This activity should be performed. Features are normally selected based on results of traffic analysis.
Set policies/procedures for network operation and maintenance.	Typically performed by plant engineers and maintenance personnel.	May be performed by either plant or IT/MIS personnel.
Develop network security and management strategies.	Not applicable.	This activity should be performed. Network security and management strategies should comply with corporate policies and procedures and be consistent with strategies developed for the plant/enterprise network.

4.2.3 Physical Network Design

Table 4-6 Physical Network Design

Activity	Industrial Network with Passive Infrastructure	EtherNet/IP
Select cable system components and infrastructure devices.	This activity is equivalent to selection of the network cable system components.	This activity must be performed. System integrators and IT Dept. may help.

4.2.4 Design Verification

Design verification activities (Table 4-7) should be performed for both isolated and non-isolated EtherNet/IP networks and should include testing, optimization, and documentation of the network design.

Table 4-7 Design Verification

Activity	Industrial Network with Passive Infrastructure	EtherNet/IP
Test network design.	Normally performed, usually consists of a test pilot design and implementation.	Should be performed and may consist of design and implementation of a test pilot.
Optimize network design.	Performed if necessary.	Should be performed if necessary.
Document network design.	Usually performed.	Should be performed or will be difficult to troubleshoot when key personnel are no longer available.

A more detailed description of these activities is provided in the following chapters.

4.2.5 Summary of Key Points on EtherNet/IP Design

- EtherNet/IP networks use an active infrastructure; they need routers and switches.
- EtherNet/IP is used for control applications; subnets can handle real-time applications.
- Control traffic on EtherNet/IP is based on the producer-consumer model, resulting in IP multicast of critical control data.
- There are two types of EtherNet/IP networks: isolated and non-isolated.
- For isolated networks, the network design process is similar to device-level networks. Non-isolated networks require additional design considerations. Non-isolated networks interface to enterprise and Wide Area Networks (WANs) and the Internet to provide

information wherever and whenever it is needed by qualified users, such as plant and maintenance engineers or managers.

4.3 EtherNet/IP Infrastructure Design Issues

4.3.1 Full-Duplex vs. Half-Duplex Implications

Ethernet is based on Carrier Sense Multiple Access/Collision Detect (CSMA/CD) technology. In its original design, this technology put all nodes on a common circuit, and they could all "talk" at the same time, which created a problem—collisions of data—because when nodes talked they could not simultaneously listen. It was analogous to two-way radios with push-to-talk switches. If two radio operators simultaneously responded to the same caller, the caller heard nothing but garble from the two radios transmitting at the same time and interfering with one another. Neither operator had any way of knowing that the other was talking. Thus, Ethernet nodes had to deal with how collisions were handled and how connected devices could maximize their ability to transmit Ethernet without monopolizing the time such that other nodes could not transmit.

The issue was that Ethernet was based on half-duplex (HDX) mode of operation for the connected devices as shown in Figure 4-2. Thus, a node could only transmit data *or* receive data. It could not do both at the same time. Half-duplex mode is similar to the traffic problem created when road construction causes one lane of a two-lane highway to be closed. Traffic from both directions is trying to use the same lane. This means that traffic coming from one direction must wait until traffic coming from other direction stops.

Ethernet Talks and Listens at the Same Time

With advances in switch technology, Ethernet connections have moved toward full-duplex operation (FDX), enabling connected devices to both send and receive packets of Ethernet data at the same time.

Ethernet Half-Duplex (HDX) and Full-Duplex (FDX)

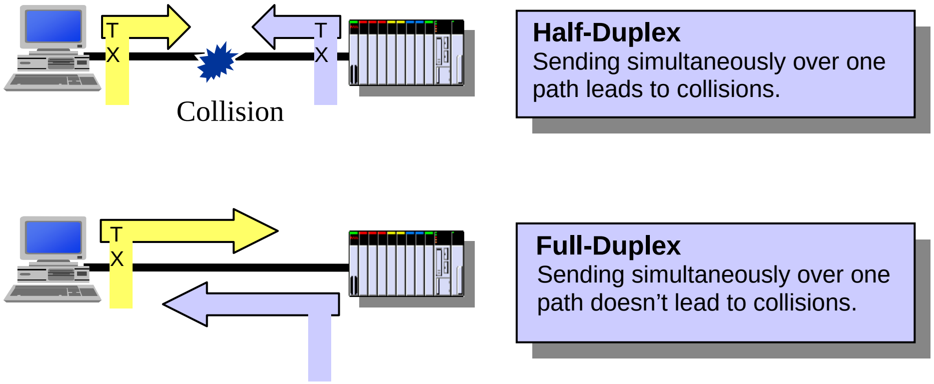
TX
TX
Collision
Half-Duplex
Sending simultaneously over one path leads to collisions.
TX
Full-Duplex
Sending simultaneously over one path doesn't lead to collisions.

Figure 4-2 Full-Duplex Operation Eliminates Collisions.

All of the Ethernet data communications equipment can be connected to additional switches and routers to extend the overall physical network size. The use of fiber-optic cabling along with twisted pair helps extend the Ethernet network. Switches and routers allow full-duplex operation—they transmit over a dedicated transmit pair of wires (or channel) and receive over a dedicated receive channel. With full-duplex, fully switched installations, collisions will be eliminated, allowing Ethernet to be applied to real-time control applications.

4.4 Layout and Wiring Considerations

There are no rigid guidelines for deciding which type of cable to use or how to route it. However, the following guidelines can make layout and wiring decisions easier by helping users compare the benefits and limitations of various options to their specific application needs. First, Figure 4-3 provides a look at a typical design.

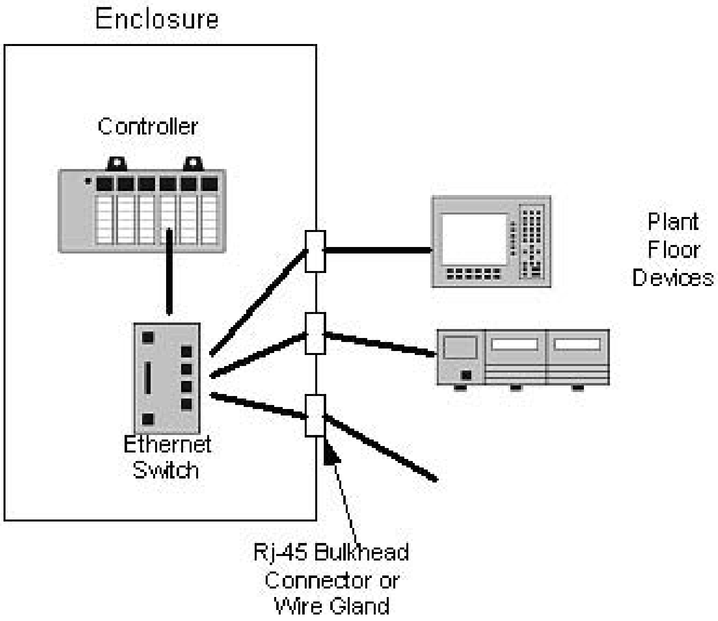

Figure 4-3 Layout and Wiring of a Typical System.

4.4.1 Length of Run and Category of Cable

In general, the total length of each cable channel (segment) in an Ethernet system must not exceed 100 meters (328 ft) between any two active devices. (See Table 4-8.) For more information consult the *EtherNet/IP™ Media Planning and Installation Guide*.[16]

[16] ODVA, Inc., *EtherNet/IP™ Media Planning and Installation Guide,* Ann Arbor: 2006. www.odva.org

Table 4-8 Maximum Cable Lengths by Type

Cable	Maximum Segment Length
10Base-T	100 meters (328 ft)
100Base-TX	100 meters (328 ft)
10Base-FL	2000 meters (6562 ft)
100Base-FX (62.5 micron multimode glass fiber)	412 meters (half-duplex) (1352 ft) 2000 meters (full-duplex) (6562 ft)
100Base-FX (single-mode fiber)	412 meters (half-duplex) (1352 ft) 15 km or more (full-duplex) (6.2 mi) Consult vendor for cable-specific limits.
1000Base-TX	100 meters (328 ft)
1000Base-SX (62.5 micron multimode glass fiber)	550 meters (1787 ft)
1000Base-LX (single-mode fiber)	5 km (3 mi) up to 10 km (6 mi) based on cable type Consult vendor for cable-specific limits.
1000Base-LH (single-mode fiber)	Up to 120 km (75 mi) based on cable type and signal frequency (1310 nm vs. 1550 nm) Consult vendor for cable-specific limits.

These are maximum suggested cable lengths. For more specific Ethernet wiring rules, refer to ANSI/TIA/EIA Std 568-B.1.[17] For fast Ethernet, low noise, high-performance 24 AWG cables and connectors are highly recommended. Any degradation in noise rejection will degrade an entire link between a switching device and the peripheral equipment.

For Gigabit operation, the increase in speed also increases the potential for data corruption by electrical noise, and requires additional attention to the cabling layout of the installation. Cat6 data cable is recommended for use with Gigabit Ethernet for higher noise immunity. Be aware of distance constraints for multimode fiber as it is typically no more than 550 meters (about 600 yds) for Gigabit Ethernet. This is compared to 2 km (1 mi) or more for 100-Mbps Ethernet.

[17] ANSI/TIA/EIA-568-B Series, Commercial Building Telecommunications Cabling – Part 1: General Requirements. http://www.tiaonline.org/

4.4.2 Cable Routing

A proper cable routing plan is essential. EtherNet/IP or any network cables should not be routed near equipment that generates strong electric or magnetic fields. In particular, routing network cables near and around the following should be avoided:

- Lights, especially those using ballasts such as mercury vapor, sodium, or fluorescent
- Motors—any size, especially if their speeds are controlled by variable frequency drives
- Drive controllers
- Arc welders
- Conduit carrying high-voltage/high currents

4.4.3 Categories of Conductors/Wiring

There are three categories of conductors for wiring in any manufacturing facility.

Category One consists of

- ac power lines,
- high-power digital ac I/O,
- high-power digital dc I/O, and
- power connections (conductors) from motor drives to motors.

Category Two consists of

- analog I/O lines,
- dc power lines for analog devices,
- low-power digital ac/dc I/O lines,
- low-power digital I/O lines, and
- EtherNet/IP.

Category Three typically includes

- low-voltage dc power lines and
- communication cables between system components within the same enclosure.

4.4.4 General Design Wiring Guidelines

The following guidelines should be followed for wiring all EtherNet/IP cables:

- If an EtherNet/IP cable must cross power lines, it should do so at right angles.
- EtherNet/IP cabling should be routed at least 1.5 meters (5 ft.) away from high-voltage enclosures or sources of radio frequency (RF)/microwave radiation.
- If the EtherNet/IP cable is running through a metal wire-way or conduit, each section of the wire-way or conduit must be bonded to each adjacent section so that it has electrical

continuity along its entire length. It also must be bonded to the enclosure at the entry point.

For more information on general wiring guidelines, see *Industrial Automation Wiring and Grounding Guidelines*[18] and the Telecommunications Industry Association publication ANSI/TIA/EIA-607, *Grounding and Bonding Requirements.*[19]

4.4.5 Wiring External to Enclosures

Cables that run outside protective enclosures may be relatively long. To minimize cross-talk from nearby power lines, it is good practice to maintain maximum separation between the Ethernet cable and other potential noise conductors. The following guidelines should be followed for routing network cables near potential noise-inducing cables.

Within a contiguous metal wire-way or conduit

- Category 1 conductors of less than 20 amps: maintain a 0.08-meter (3-in.) separation distance.
- AC power lines of 20 amps or more, up to 100 KVA: maintain a 0.15-meter (8-in.) separation distance.
- AC power lines greater than 100 KVA: maintain a 0.3-meter (12-in.) separation distance.

In open-air cabling systems

- Category 1 conductors of less than 20 amps: maintain a 0.15-meter (8-in.) separation distance.
- AC power lines of 20 amps or more, up to 100 KVA: maintain a 0.3-meter (12-in.) separation distance.
- AC power lines greater than 100 KVA: maintain a 0.6-meter (24-in.) separation distance.

Local, regional, and national codes regarding the grouping of cables should also be followed. In the absence of these codes, the general rule for noise protection is a minimum distance of 7.6 cm (3 in.) from electric light and power conductors and an additional 2.5 cm (1 in.) for each 100 volts over the 100-volt level. See Table 4-9 for outside wiring recommendations.

[18] Rockwell Automation, *Industrial Automation Wiring and Grounding Guidelines.* Publication 1770-4.1. Milwaukee: Rockwell Automation, 1998. http://literature.rockwellautomation.com/

[19] TIA-J-Std 607-A, Commercial Building Grounding (Earthing) and Bonding Requirements for Telecommunications. http://retail.ihs.com/abstracts/tia/tia-j-std-607a.jsp

Table 4-9 Outside Wiring

Voltage Level	Minimum Distance
0-100 V	7.6 cm (3 in.)
101-200 V	10.2 cm (4 in.)
201-300 V	12.7 cm (5 in.)
301-400 V	15.2 cm (6 in.)
401-500 V	17.8 cm (7 in.)

4.4.6 Wiring Inside Enclosures

Cable sections that run inside protective equipment enclosures are relatively short. As with wiring external to enclosures, maximum separation between Ethernet cables and Category-1 conductors should be maintained.

When running cable inside an enclosure, conductors should be routed external to all raceways in the same enclosure or in a raceway separate from Category-1 conductors. Cable should be routed at least the distance from the noise source listed in Table 4-10.

Table 4-10 Routing Cable Around Various Noise Sources

Minimum Distance	Strength of Noise Source
0.08 meter (3 in.)	Category-1 conductors of less than 20 amps
0.15 meter (8 in.)	AC power lines of 20 amps or more, up to 100 KVA
0.6 meter (24 in.)	AC power lines greater than 100 KVA

4.4.7 Cable Grounding Considerations

For information regarding grounding and bonding requirements for your network, the appropriate local, regional, and national or international codes should be followed.

Grounding of Shielded Cables

Shields play an important role in providing noise immunity for systems. However, an improperly installed shielded cable can cause problems due to voltage offsets in a grounding system. To minimize the effects of ground offsets, the shield should be isolated at one end of the cable. In this case, the shield should be isolated at the device. The ground can be applied at the switch or other infrastructure component. It should be noted that some European wiring practices recommend grounding the cable shields at both ends, requiring that particular care be taken to avoid ground loops. The quality of a plant's grounding system must be thoroughly evaluated prior to using this approach.

Eliminating ground loops is extremely important in reducing noise interference caused by ground offsets or local ground transients. Ground offsets occur when there is a potential difference between two earth-ground points in a system as shown in Figure 4-4. This potential difference can consist of continuous dc or ac voltages as well as transients.

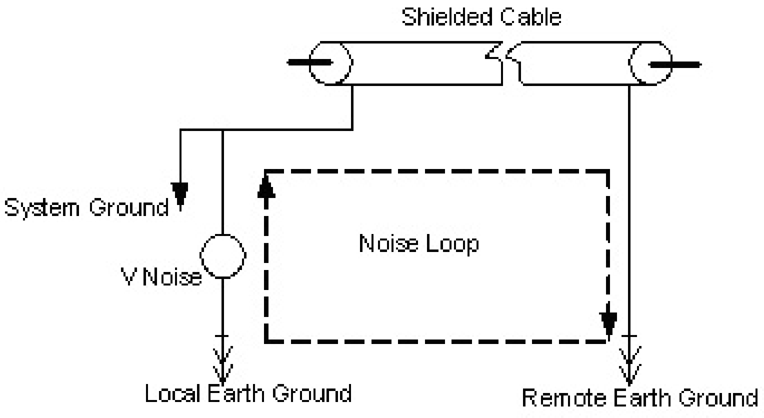

Figure 4-4 Ground Noise Loop in Shielded Cable.

A cable that connects the two points together provides a secondary travel path for current. The current through the communications cable's shield will couple noise into the communications system. This noise will have a direct impact on the signal-to-noise performance of the system.

There can be as much as a 45-volt high frequency offset in ground potential between the two ends of a 100-meter (328 ft) cable. This offset can cause noise currents in the shields. Consequently, to eliminate ground loops, the cable should be grounded at the switch end only as shown in Figure 4-5.

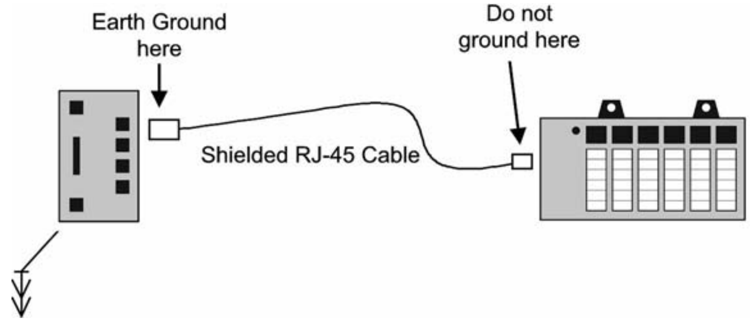

Figure 4-5 Grounding of Cable Shield.

If, for example, a device other than a switch provides a low resistance (<500k ohms) ground at the jack, the shield should not be connected at the device end of the cable. Rather, the shield should be cut back and insulated from the connector plug’s shell to break the ground as shown in Figure 4-6.

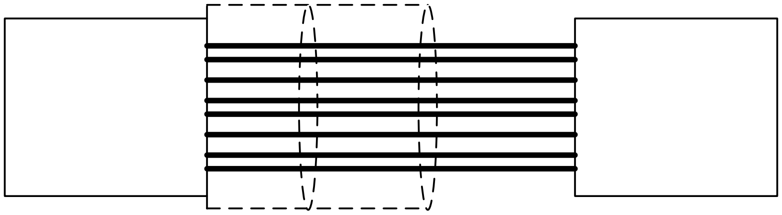

Figure 4-6 Null Termination of Shielded Cable.

4.5 Industrial Ethernet Connectors

4.5.1 Attributes of Industrial Ethernet Cables

When selecting network cable, the jacket construction must be compatible with any vibration and the temperature and chemicals in the environment. It is important not to overlook cable electrical specifications, such as over-temperature, as many off-the-shelf cables do not meet the TIA/EIA standards at industrial temperatures. Cable jackets may by easily damaged, even at low temperature ranges. Chemicals can be absorbed into the jackets and wire insulation, causing plastic deterioration and performance degradation. Sealed media may be required if connectors or cable jackets are exposed to a harsh environment.

In industrial unshielded cable installations, the use of balanced cables is recommended. If the networking application is in a high-noise environment, then shielded cable media or fiber-optic cabling should be considered.

4.5.1.1 Oil-resistant Jackets

If an application requires control of equipment that uses cutting oils or lubricating chemicals, cables constructed with oil-resistant jackets should be specified.

4.5.1.2 Plenum-rated Cables

Plenum installations require special material compounds in the makeup of cables. If an application requires cables to be run in air-handling space, plenum-rated cables should be used, and may be required by local electrical codes. (Often local electrical codes are concerned with voltage and current levels for power wiring in a heated duct.) Another cautionary note: Many air-handling ducts have sharp edges of sheet metal throughout—especially at joints and corners, which can easily cut cable when it is pulled through existing duct work.

4.5.1.3 Weld Splatter-Proof Cables

If an application requires control of welding equipment, cables should be routed to reduce damage from weld splatter and noise ingress. All cables should be routed to cross over welding and motor control cables at right angles and should never run parallel to control cables. The cables should be protected from weld splatter either by adding protective sheath, using conduit, or selecting cables with the proper jacket insulation.

4.5.1.4 High Flex

Typically, Ethernet cables are constructed of 24-gauge solid copper conductors, which are not suited to constantly moving machinery, such as robots. High-flex applications require stranded-type conductor construction to extend the life of a cable in flexing applications. Also note that high-flex cables are rated differently based on the type of flexing to which they will be exposed. Trailing cable or “C Track” applications require different cables than applications where cables are exposed to a continuous back and forth, or “tick-tock” type of flexing. Other technologies that can replace wiring in these applications include wireless and infra-red (IR) communications.

4.5.1.5 Ambient Temperature

The ambient temperature of the environment must be considered in specifying the cables and connectors to be installed. Some Ethernet cables cannot survive or will not perform to specification in the extreme hot and cold temperatures of some industrial environments.

4.6 Performance Considerations: The Need for Switches

Managing network traffic involves knowing what hardware to apply, such as a switch versus a router. It also requires logical segmentation of the network. Using Virtual LANs (VLANs), for instance, can be compared to assigning a lane on a highway for trucks only. Finally, just as a freeway system gives priority to emergency vehicles, network traffic must be prioritized. The objective is to use various mechanisms to assure that critical messaging traversing the network will always have enough resources to guarantee it will arrive—and arrive on time.

4.6.1 Performance of Switches

In the past, nodes were simply connected together using hubs wherever possible to keep costs low. Their limitations have already been covered. While hubs provide an easy way to add additional nodes and extend the length of the network, they do not break up the actual network into discrete segments. That is a job for switches.

Switches are a fundamental part of most networks. They make it possible for several users to send information over a network at the same time without slowing each other down.

In a fully switched network, switches replace all the hubs of an Ethernet network with a dedicated segment for every node. Since the only devices on each segment are the switch and the node, the switch picks up every transmission before it reaches another node. The switch then forwards the frame over the appropriate segment. Since any segment contains only a single node, the frame only reaches the intended recipient. Before switches were used, Ethernet was a half-duplex network. Thanks to managed and unmanaged switches, Ethernet can be a full-duplex network.

Switched Networks and Full-Duplex Ethernet

In a switched network, nodes only communicate with the switch and never directly with each other. With a dedicated segment for every node, switching allows a network to maintain full-duplex Ethernet.

Switched networks employ either twisted-pair or fiber-optic cabling, both of which use separate conductors for sending and receiving data. In this environment, Ethernet nodes can forgo the collision detection process and transmit at will, since they are the only potential devices that can access the medium.

In other words, traffic flowing in each direction has a lane to itself. This allows nodes to transmit to the switch at the same time the switch transmits to them, achieving a collision-free environment. Transmitting in both directions also can effectively double the apparent bandwidth of the network when two nodes are exchanging information. For example, if the speed of the network is 100 Mbps, then each node can transmit at 100 Mbps at the same time.

4.7 Switching Technologies

Switches usually work at Layer 2 (data link) of the OSI Reference Model using MAC addresses while routers work at Layer 3 (network) with Layer 3 addresses (IP, IPX®, or AppleTalk® depending on what Layer 3 protocols are being used). The algorithm that switches use to decide how to forward packets is different from the algorithms that routers use to forward packets.

One of the differences between the algorithms of switches and routers is how each algorithm handles broadcasts. On any network, the concept of a broadcast packet is vital to the operability of the network. Whenever a device needs to send out information but does not know where to send it, it sends out a broadcast. For example, every time a new computer or other device comes onto the network, it sends out a broadcast packet to announce its presence. Broadcasts are transmitted any time a device needs to make an announcement to the rest of the network or is not sure where to send information.

A switch will pass along any broadcast packets it receives to all other segments in the broadcast domain, but a router will not. Without the specific address of another device for which a message is intended, a router will not let the data packet pass through. This is an effective way of keeping networks separate from each other but is not helpful when different parts of the same network need to communicate with one another. Switches provide a solution to this problem.

Ethernet Switches Rely on Packet-Switching.

The switch establishes a connection between two segments and maintains the connection just long enough to send the current packet. Incoming packets, which are part of an Ethernet frame, are saved to a temporary memory area called a buffer. The MAC address contained in the frame's header is read and then compared to a list of addresses maintained in the switch's lookup table.

Today's packet-based switches normally use the store-and-forward method of routing traffic. Previously installed switches may also have used the cut-through, or modified cut-through traffic routing techniques.

The cut-through method of switching is based on the premise that there is no need for the switch to wait for the arrival of the complete frame. The switch waits only long enough to read the destination address in the frame before it begins forwarding the frame to its destination.

A switch using store-and-forward will save the entire packet to the buffer and run a Cyclic Redundancy Check (CRC) to locate transmission and storage errors or other problems. If the packet has an error, then it is discarded. Otherwise, the switch looks up the MAC address and sends the packet on to the destination node. Some switches combine the two methods by using cut-through until a certain error level is reached, then changing over to store-and-forward. Very few switches are strictly cut-through since this provides no error detection.

Modified cut-through (also known as fragment-free) switches are an attempt to offer the best of both store-and-forward and cut-through switching. In an Ethernet environment, an incoming frame is held until the first 64 bytes have been received. If the frame is incomplete or corrupt, it can usually be detected within the first 64 bytes. Therefore, a trade-off between switch latency and error checking is achieved.

Message Types

There are three types of communication in IP networks: unicast, multicast and broadcast. A unicast packet is addressed to an individual node, a multicast packet is addressed to a group of nodes, and a broadcast packet is addressed to all nodes on a network. Each node, regardless of the type of communication in which it participates, must have a uniquely assigned IP address. When another node wants to send a unicast packet to this node, it must use this address as a destination address. When a node wants to send to this node, or a group of nodes, it must use a multicast packet and a destination address, a specific address selected from the IP multicast address class (class D). When a node wants to send a broadcast packet addressed to all nodes on the network, it must use, as a destination address, IP address 255.255.255.255.

4.8 IGMP Snooping

Internet Group Management Protocol (IGMP) snooping constrains the flooding of multicast traffic by dynamically configuring the switch interfaces such that multicast traffic is forwarded only to those interfaces associated with IP multicast devices. In other words, when a multicast message is sent to the switch, the switch forwards the message only to the interfaces that are interested in this traffic.

IGMP Snooping

The Internet Group Management Protocol is very important because it reduces the load of traffic traversing through the network. It also relieves the hosts from processing frames that are not needed.

In the producer-consumer model used by EtherNet/IP, IGMP snooping limits unnecessary traffic from the producing I/O, such that it only reaches the devices that consume that data. Messages delivered to a given device (say Device A) that were intended for another device (say Device B) will consume resources, negatively affecting the performance of Device A.

Networks with multiple multicasting devices will suffer performance issues if IGMP snooping or other multicast limiting schemes are not implemented. The number of devices that can be supported in a given network varies with the type of traffic that is present.

When using IGMP snooping, at least one IGMP querier is needed somewhere on the network. A router or switch can act as the IGMP querier if the device supports this functionality. An IGMP querier periodically queries the multicast receiving devices as to their multicast interest. Once queried, a multicast recipient will respond with the multicasts it wishes to receive by sending an IGMP Join Report for each multicast group in which it has interest.

The Ethernet switch snoops on the IGMP traffic coming to that switch and keeps track of multicast groups and member ports. When the switch receives an IGMP Join Report from a host for a particular multicast group, the switch adds the host port number to the associated multicast forwarding table entry. When it receives an IGMP Leave Group message from a host, it removes the host port from the table entry. After it relays the IGMP queries, it deletes entries periodically if it does not receive any IGMP membership reports from the multicast clients. A Layer 3 device (router) normally performs the querying function, but it can also be implemented by another network element designed to take over this functionality.

When IGMP snooping is enabled in a network with Layer 3 devices, the multicast router sends out periodic IGMP general queries to all VLANs. The switch responds to the router queries with only one join request per MAC multicast group, and the switch creates one entry per VLAN in the Layer 2 forwarding table for each MAC group from which it receives an IGMP join request. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry.

Layer 2 multicast groups learned through IGMP snooping are dynamic. However, in a managed switch, it is possible to statically configure MAC multicast groups. If a group membership for a multicast group address is specified statically, this setting supersedes any automatic manipulation by IGMP snooping. Multicast Group Membership Lists can consist of both user-defined and IGMP snooping-learned settings.

If a port spanning-tree, a port group, or a VLAN ID change occurs, the IGMP snooping-learned multicast groups from this port on the VLAN are deleted.

4.8.1 Joining a Multicast Group

When a host connected to the switch wants to join an IP multicast group, it sends an IGMP join request, specifying the IP multicast group it wants to join. When the switch receives this message, it adds the port to the IP multicast group port address entry in the forwarding table.

In Figure 4-7, Host 1 wants to join multicast group 224.1.2.3 and multicasts an unsolicited IGMP Join Report (IGMP join message) to the group with the equivalent MAC destination address of 0100.5E01.0203. The switch recognizes IGMP packets and forwards them to the CPU.

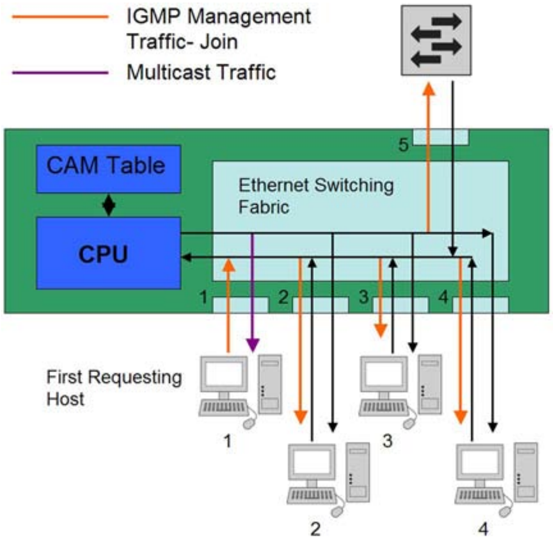

Figure 4-7 Initial IGMP Join Message.

When the CPU receives the IGMP report multicast by Host 1, the CPU uses the information to set up a multicast forwarding table entry as shown in Table 4-11. This entry includes the port numbers of Host 1 and the router.

Table 4-11 Multicast Forwarding Table Entry

IP Multicast Forwarding Table Destination Address	Type of Packet	Ports
0100.5E01.0203	!IGMP	1, 2

The switch architecture allows the CPU to distinguish IGMP information packets from other packets for the multicast group. The switch recognizes the IGMP packets through its filter engine. This prevents the CPU from becoming overloaded with multicast frames.

The entry in the multicast-forwarding table tells the switching engine to send frames addressed to the 0100.5E01.0203 multicast MAC address that are not IGMP packets (!IGMP) to the router and to the host that has joined the group.

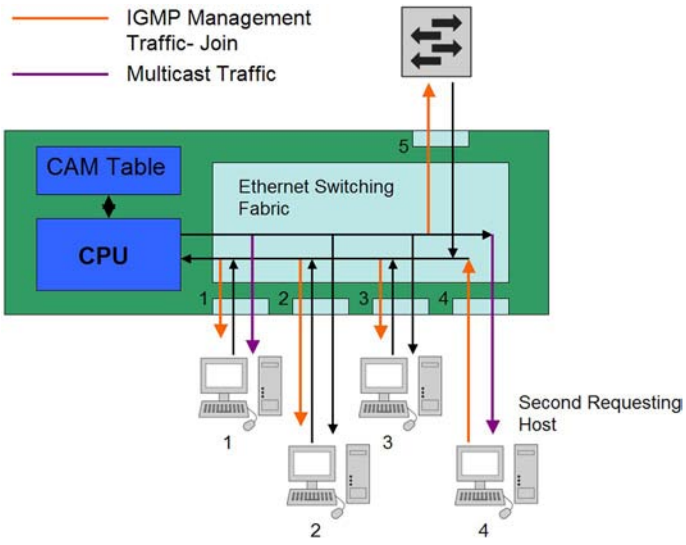

Figure 4-8 Second Host Joining a Multicast Group.

If another host (for example, Host 4 in Figure 4-8) sends an IGMP join message for the same group, the CPU receives that message and adds the port number of Host 4 to the multicast forwarding table as shown in Table 4-12.

Table 4-12 Updated Multicast Forwarding Table Entry

Updated Multicast Forwarding Table Destination Address	Type of Packet	Ports
0100.5E01.0203	!IGMP	1, 2, 5

4.8.2 Leaving a Multicast Group

The router sends periodic IP multicast general queries, and the switch responds to these queries with one join response per MAC multicast group. As long as at least one host in the VLAN needs multicast traffic, the switch responds to the router queries, and the router continues to forward the multicast traffic to the VLAN. The switch only forwards IP multicast group traffic to those hosts listed in the forwarding table for that IP multicast group.

When hosts need to leave a multicast group, they can either ignore the periodic general-query requests sent by the router, or they can send a leave message. When the switch receives a leave message from a host, it sends out a group-specific query to determine if any devices behind that interface are interested in traffic for the specific multicast group. If, after a number of queries, the router processor receives no reports from a VLAN, it removes the group for the VLAN from its multicast-forwarding table.

4.8.3 Immediate-Leave Processing

IGMP snooping immediate-leave processing allows the switch to remove an interface that sends a leave message from the forwarding table without first sending out MAC-based general queries to the interface. The VLAN interface is pruned from the multicast tree for the multicast group specified in the original leave message. Immediate-leave processing ensures optimal bandwidth management for all hosts on a switched network, even when multiple multicast groups are in use simultaneously.

4.9 Quality of Service in a Switched Network

Phrases like "gigabits of back-plane capacity," "millions of switched packets per second," and "non-blocking switch fabrics" reflect the high-performance of today's Ethernet switches and typically generate a simple question: Why the need for Quality of Service (QoS)? The answer is congestion.

A switch may be the fastest switch in the world, but if either of the two scenarios in Figure 4-9 is present, the switch will experience congestion.

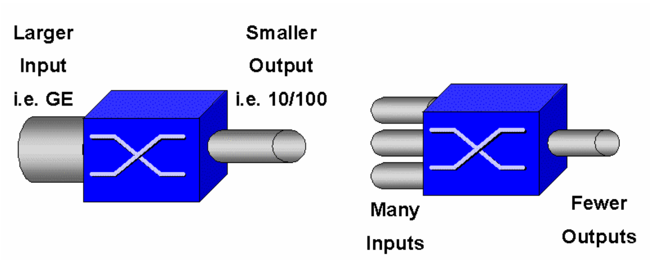

Figure 4-9 Inputs Higher in Speed or Quantity than Outputs Cause Congestion.

If the congestion management features on the switch are not up to par during these congested periods, performance will suffer and packets will be dropped. In a TCP/IP network, packet drops will generate retransmissions, which then increases network load. In networks that are already congested, this increase in network load further exacerbates existing network performance issues.

Latency-sensitive traffic, such as motion control messages, can be severely affected if transmission delays occur. Adding more buffers to a switch will not necessarily alleviate congestion problems since latency-sensitive traffic needs to be switched as quickly as possible.

To address network congestion issues, different stages of QoS need to be implemented:

- First, identify different traffic types in the network using classification techniques.
- Next, implement advanced buffer management techniques to avoid high-priority traffic from being dropped during congestion.
- Finally, incorporate scheduling techniques to transmit high-priority traffic from queues as quickly as possible.

4.9.1 QoS Flow

The common definition of QoS in Layer 2 switches is to prioritize native encapsulated Ethernet frames or to honor 802.1p Class of Service (CoS) tagged Ethernet frames.[20] However, advanced QoS mechanisms take this definition a step further. Layer 2 QoS on advanced Ethernet switches entail up to four distinct stages: classification, policing, marking, and queue/schedule. As shown in Figure 4-10, the first three represent actions taken at ingress, and queue/schedule is an action taken at egress.

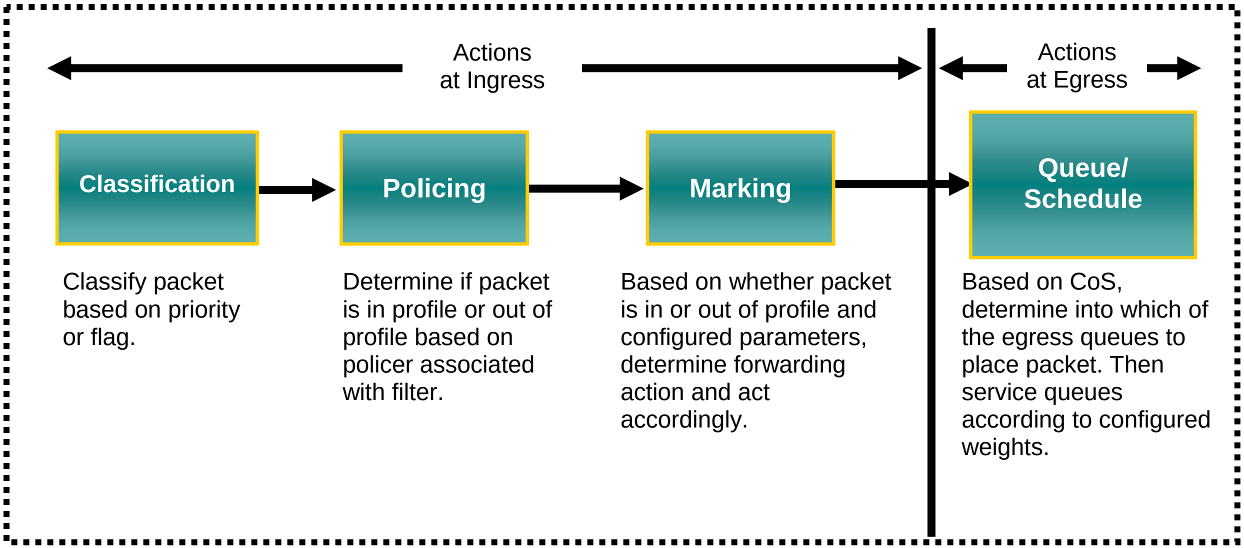
Actions at Ingress
Actions at Egress
Classification
Policing
Marking
Queue/ Schedule
Classify packet based on priority or flag.
Determine if packet is in profile or out of profile based on policer associated with filter.
Based on whether packet is in or out of profile and configured parameters, determine forwarding action and act accordingly.
Based on CoS, determine into which of the egress queues to place packet. Then service queues according to configured weights.

Figure 4-10 Quality of Service Stages.

Feature-rich switches can be used to study and interpret the flow of QoS traffic through the switch from ingress, as it gets processed through the switch fabric, and as it flows out (egress).

A switch can be configured to prioritize frames based on given criteria at different layers of the OSI Model's stack. For example, priorities may be assigned according to the source MAC Address (Layer 2) or the destination TCP port (Layer 4). Any traffic traveling through the interface to which these criteria apply will be classified and tagged within the frame header, assigning them a given priority. Once the packet has been classified, it can be policed/marked down (if necessary), and it will be queued according to the specified priority. Once it has been placed into the appropriate holding queue for transmission on egress, it will be scheduled based on the configuration of the scheduling algorithm.

4.9.2 Class of Service

Class of Service (CoS) refers to three bits on an 802.1Q header that are used to indicate the priority of the Ethernet frame as it passes through a switch network. The CoS bits in the 802.1Q header are commonly referred to as the 802.1p bits. Not surprisingly, there are three CoS bits that match the number of bits used for IP Precedence (which is part of the Type of Service [ToS] byte). In many networks, to maintain end-to-end QoS, a packet may traverse both Layer 2 and Layer 3 domains, thus the ToS and CoS fields should be mapped to each other.

[20] IEEE Std 802.1p is a Standard used in IEEE Std 802.1D, Standard for Local and Metropolitan Area Networks: Media Access Control Bridges, and used in IEEE Std 802.1Q, IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks. http://standards.ieee.org/getieee802/802.1.html

Figure 4-11 depicts an Ethernet frame tagged with an 802.1Q field, which consists of a two-byte Ethertype and a two-byte 802.1Q Tag. Within the two-byte tag are the user priority bits (known as 802.1p CoS bits).

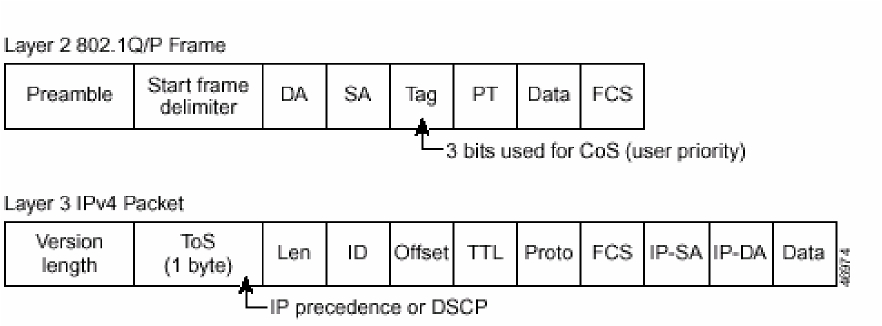

Figure 4-11 Details of Ethernet Frames and IP Packet.

4.9.3 Priority Mechanisms

For any QoS services to be applied to data, there must be a way to "tag" or prioritize an IP packet or an Ethernet frame. The Layer 2 802.1p Class of Service (CoS) and Layer 3 IP Type of Service (ToS) are the fields used to achieve this as shown in Figure 4-12.

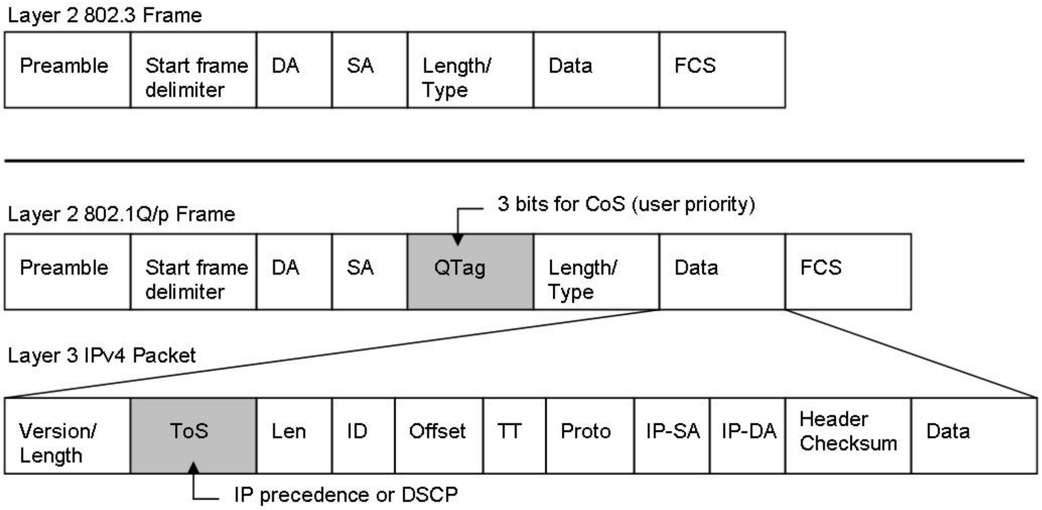

Figure 4-12 Prioritization of IP Packets.

There are several parameters to take into account when considering Layer 2 through Layer 4 classification. A switch can implement Layer 2 classification (as shown in the "Encapsulated Packet" diagram within Figure 4-12) by looking at the Layer 2 Header and classifying the frame accordingly. An advanced managed switch may also view any CoS value that comes tagged on a frame coming into the switch and will use this information to determine the priority that it should be assigned, then take the appropriate actions. Feature rich switches will also process the packets based on Layer 3 IP address information as well as Layer 4 (TCP and UDP) information, and the Differentiated Services Code Point (DSCP) value.

4.10 Virtual Local Area Networks

As networks have grown in size and complexity, many companies have turned to Virtual Local Area Networks (VLANs) to provide a means for structuring this growth logically. Basically, a VLAN is a network that is created logically, using software, as opposed to physical cabling. Therefore, only devices predefined according to a specific criterion for the same broadcast domain will receive a transmission generated by a station on a VLAN.

A broadcast domain is a network (or portion of a network) that will receive a broadcast packet from any node located within that network. Since routers do not pass along broadcast messages, in a typical network everything on the same side of the router is part of the same broadcast domain.

Switches Implemented with VLANs

A switch implemented with VLANs has multiple broadcast domains similar to a router. The switching device used to construct a VLAN can be an advanced managed Layer 2 switch that operates at the MAC sub-layer of the OSI Model, or a Layer 3 switch that operates at the network layer.

Here are some common reasons why a company might want to install VLANs:

- Security. Separating systems with sensitive data from the rest of the network decreases the chance that someone will gain access to information they are not authorized to see.
- Projects/special applications. Managing a project or working with a specialized application can be simplified by using a VLAN to bring all the required nodes together.
- Performance/bandwidth. Monitoring of network use allows the network administrator to create VLANs that reduce the number of router hops and increase the apparent bandwidth for network users.
- Broadcasts/traffic flow. Since a principle element of a VLAN is that it does not pass broadcast traffic to nodes that are not part of the VLAN, it automatically reduces broadcasts. Access lists provide the network administrator with a way to control who sees what network traffic.
- Departments/specific job types. Companies may want to set up VLANs for departments that are heavy network users (such as Multimedia or Engineering). A VLAN dedicated to specific types of employees from different departments (such as managers or sales people) might also be set up across departments.

While more than one VLAN can exist on a single switch, these VLANs cannot communicate directly with each other. If they could, it would defeat the purpose of having a VLAN, which is to isolate a part of the network. Communication between VLANs requires the use of a router.

4.10.1 VLAN Trunking

VLANs can span across multiple switches and you can have more than one VLAN on each switch. For multiple VLANs on multiple switches to be able to communicate via a single link between the switches, trunking is needed. Trunking is the technology that allows information from multiple VLANs to be carried over just one link between switches.

Switches use VLAN trunking protocols to communicate VLAN configuration information between them. Standardized protocols such as the Group VLAN Registration Protocol (GVRP) allow VLAN trunk connections between switches from different vendors. Proprietary, vendor-specific trunking protocols are also available.

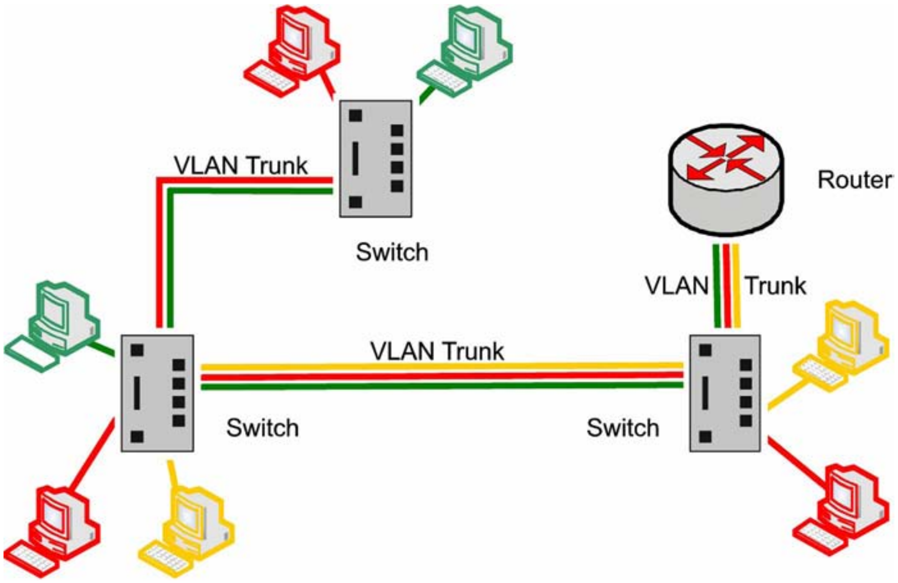

Figure 4-13 Devices on Three VLANs Communicate via Trunking Links between Three Switches and a Router.

In Figure 4-13, each switch has two VLANs. On the first switch, VLAN A and VLAN B are sent through a single port (trunked) to both the router and through another port to the second switch. VLAN C and VLAN D are trunked from the second switch to the first switch and through it to the router. This trunk can carry traffic from all four VLANs. The trunk link from the first switch to the router can also carry all four VLANs. In fact, this one connection to the router actually allows the router to appear on all four VLANs as if it had four different physical ports connected to the switch.

The VLANs can communicate with each other via the trunking connection between the two switches using the router. For example, data from a computer on VLAN A that needs to get to a computer on VLAN B (or VLAN C or VLAN D) must travel from the switch to the router and back again to the switch. Because of the switches' transparent bridging algorithm and trunking, both PCs and the router think they are on the same physical segment.

4.11 Multilayer Ethernet Switches

A multilayer Ethernet switch performs the traditional function of a Layer 2 device but also can perform functions based on information from other layers of the networking stack. For example, a multilayer switch can inspect the IP information encapsulated in an Ethernet frame and perform routing decisions based on it (Layer 3 switching), assign a given priority to the frame (QoS), and deny or permit access (security). A Layer 4-aware switch could do the same based on UDP or TCP ports (source or destination). This can be a very useful feature in EtherNet/IP networks as a switch can be configured to give high priority to the Common Industrial Protocol (CIP) traffic based on the UDP port destination/source in the datagram. It also could prevent unauthorized users from accessing the network by checking their IP addresses.

4.11.1 Layer 2 and Layer 3 Switching

Switching is the process of taking an incoming frame from one interface and delivering it out through another interface. Routers use Layer 3 switching to route a packet, and switches (Layer 2 switches) use Layer 2 switching to forward frames.

The Difference between Layer 2 and Layer 3 Switching

The difference between Layer 2 and Layer 3 switching is the type of information inside the frame that is used to determine the correct output interface. With Layer 2 switching, frames are switched based on MAC address information. With Layer 3 switching, frames are switched based on network-layer information.

Traditional Layer 2 switching does not look inside a packet for network-layer information as does Layer 3 switching. Layer 2 switching is performed by looking at a destination MAC address within a frame. It looks at the frame's destination address and sends it to the appropriate interface if it knows the destination address location. Layer 2 switching builds and maintains a switching table that keeps track of which MAC addresses belong to each port or interface.

If the Layer 2 switch does not know where to send the frame, it transmits the frame out on all its ports and is said to flood the ports to learn the correct destination. When the frame's reply is returned, the switch learns the location of the new address and adds the information to the switching table.

Layer 2 Addresses

Layer 2 addresses are determined by the manufacturer of the data communications equipment. They are unique addresses that are derived in two parts: the manufacturing (MFG) code and the unique identifier. The MFG code is assigned to each vendor by the IEEE. The vendor assigns a unique identifier to each board it produces. Layer 2 addresses are fixed with a device, whereas Layer 3 addresses can be changed. In addition, Layer 2 addresses assume a flat address space with universally unique addresses.

Layer 3 switching operates at the network layer. It examines packet information and forwards packets based on their network-layer destination addresses. Layer 3 switching also supports router functionality.

As routers operate at Layer 3 of the OSI Model, they can adhere to and formulate a hierarchical addressing structure. Therefore, a routed network can tie a logical addressing structure to a physical infrastructure, for example, through TCP/IP subnet networks for each segment. Traffic flow in a switched (flat) network is, therefore, inherently different from traffic flow in a routed (hierarchical) network. Hierarchical networks offer more flexible traffic flow than flat networks because they can use the network hierarchy to determine optimal paths and contain broadcast domains.

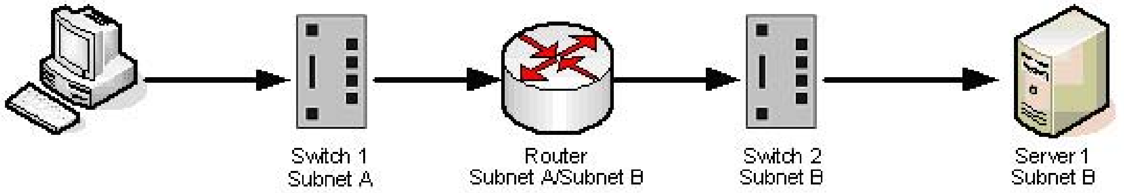

Figure 4-14 Flow of Inter-subnet Traffic with Layer 2 Switches and Routers.

As shown in Figure 4-14, for the PC client to communicate with Server 1, which is on another subnet, it must traverse through Switch 1 (a Layer 2 switch), then through the router (a Layer 3 switch) and, finally, through Switch 2 (a Layer 2 switch). There is the potential for a tremendous bottleneck, which can threaten network performance, because the inter-subnet traffic must pass from one network to another.

To relieve this bottleneck, network designers can add Layer 3 capabilities throughout the network. By implementing Layer 3 switching on edge devices, they can alleviate the burden on centralized routers.

Figure 4-15 illustrates how deploying Layer 3 switching throughout the network allows the PC client to communicate directly with Server 1 without passing through the router.

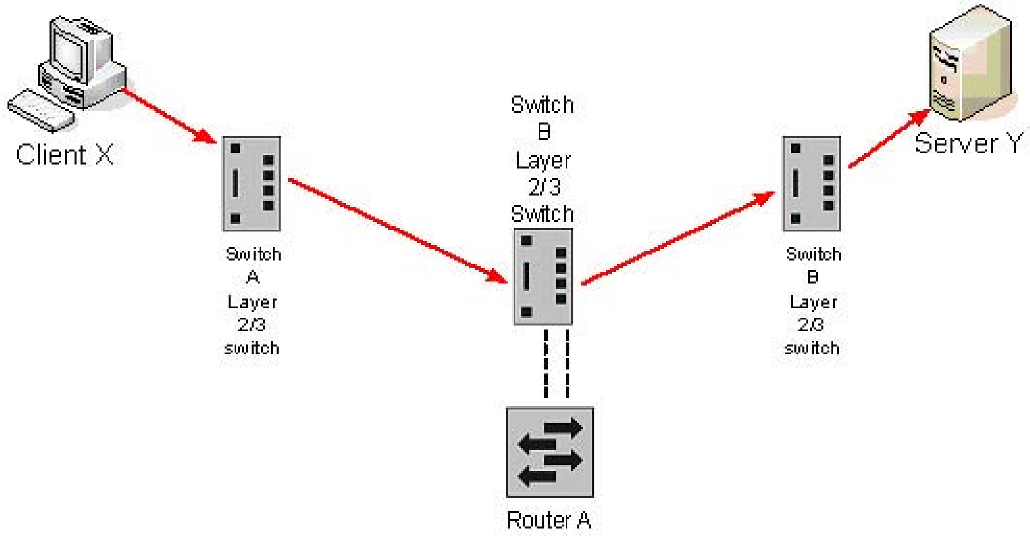
Client X
Switch B Layer 2/3 Switch
Server Y
Switch A Layer 2/3 switch
Switch B Layer 2/3 switch
Router A

Figure 4-15 Flow of Inter-subnet Traffic with Layer 3 Switches.

4.12 Achieving Deterministic Behavior

Different mechanisms can be implemented to achieve reliable and deterministic behavior on a network. Scale and criticality are two key factors for determining when to implement such mechanisms. Just as a back country road often can effectively move traffic and manage the potential for collisions with a single lane, traffic in larger metropolitan areas generally demands multi-lane, multi-level highways. As the amount of traffic increases, the need for mechanisms that can guarantee a given performance becomes evident.

In general, two factors drive the need to implement performance mechanisms in EtherNet/IP networks.

- Scale: Small networks that will not be sharing their resources with other applications, such as video, voice, non-control data, etc. have lower potential for congestion and consequent jitter, latency, and frame losses. As the networks grow and the hosts begin sharing data through a producer-consumer model, the need for mechanisms such as IGMP snooping and QoS become indispensable to guarantee a deterministic behavior.

- Criticality: Some applications can tolerate lost data without any significant impact on the manufacturing process, and other applications can be adversely affected by even small latency or packet loss. Even small disruptions can have a big economic impact on a manufacturing facility. In these cases, it is normally a requirement to implement cost-effective solutions that reduce the potential for a work stoppage.

Large EtherNet/IP deployments or deployments where a line stop is critical should consider implementing:

- A fully switched network. This will eliminate collisions and improve the deterministic behavior of the data network.

- Quality of Service (QoS) traffic prioritization. QoS prioritization allows time critical traffic to have preferential handling over supervisory traffic.

- Logical segmentation of the network. VLANs improve security and contain broadcast messaging.

- IGMP snooping. This will control multicast messages that can slow the performance of the network hosts. It also exponentially reduces the amount of traffic on the network, reducing the chance for congestion and consequent packet loss.

4.13 Managing the Interface between Control and IT

Managing the interface between control and IT requires cooperation between two disparate sets of needs. Covered here is the technology of integration—connecting EtherNet/IP networks to enterprise networks. There is no doubt that EtherNet/IP networks can solve the “islands of automation” syndrome that has plagued manufacturers for so long, but making them work with the enterprise network will require a commitment for cooperation between a manufacturer’s enterprise Information Technology Department and plant-floor network engineers.

It is important to keep in mind that everyone is working for the same team and toward the same goals. This spirit of cooperation is a prerequisite to achieving an effective plant-floor-to-enterprise information system. Business leaders also must realize that while engineers from both disciplines share the same ideals—performance, speed, and security—they tend to employ different problem-solving methods to achieve the same ends.

EtherNet/IP networks installed by control engineers will be carefully scrutinized by IT network specialists before interconnecting them to business-level applications. For example, the IT department probably will insist that any Ethernet infrastructure devices follow IT guidelines, policies, and procedures.

In addition, the IP addresses used on EtherNet/IP networks will need to be coordinated with an IT IP address administrator. Network management policies will need to be changed to give IT operations' support personnel access to Ethernet switch configuration parameters. The security threats posed by devices on the EtherNet/IP networks will require gateway and/or firewall protection for the IT network.

Conversely, control engineers are likely to insist on "owning" the EtherNet/IP network switch hardware installed on the plant floor. They will require that any piece of communication equipment conveying control signal traffic be locally monitored, readily accessible, and replaceable by on-site maintenance personnel. Control engineers may deem the IT support policies for critical production equipment to be not sufficiently real-time for the uptime demands of the factory floor, and an IT department counter-proposal to install redundant networks to ensure communication reliability may be deemed too expensive for the return-on-investment metrics of the plant. Control engineers may complain that the IP address administration policies and procedures are not workable for their plant-floor requirements, and may insist on the ability to modify or expand control system networks without outside interference.

In one respect, IT security practices may be criticized for being too excessive for the control system environment. Conversely, control system vendors are very cautious about operating system security patches, and will want to verify that any operating system patches are compatible with control systems before they are applied. Indiscriminate application of operating system and security patches has been known to shut down some control system software that has not yet been checked out for operation with the patches.

Who owns EtherNet/IP networks? This is not a simple question. EtherNet/IP networking issues span both the information technologies' and control systems' organizational boundaries.

Having looked at isolated and non-isolated networks and provided detailed instructions for designing and planning EtherNet/IP networks, as well as covering wiring issues, network topologies, and factors that affect network performance, perhaps the most important issue is the "people factor." Control engineers and IT engineers need to work together on the top-down approach presented here to build an effective plant-wide information system that gets the right data to those who need it wherever they are, and to get the data to these destinations promptly and securely.

5 Deploying the Network

EtherNet/IP networks can be tested using relatively inexpensive equipment and software. There are some quick and easy methods of assuring that EtherNet/IP cables and connectors are providing connectivity. Deployment also involves configuring the system for EtherNet/IP traffic.

Methods of self-testing routers and switches make it easy to diagnose the basic issues: Is it connected? Does it see its signal? Easy software checks can further ensure that the connections are not only made, but that hardware will communicate over TCP/IP connections. Simple hand-held testers are available, many of which offer embedded diagnostics. Using such testers, engineers can be assured that a new system will communicate without having to spend tens of thousands of dollars on test equipment.

EtherNet/IP provides a large set of simple, universal tools that can be leveraged to increase the speed of network deployments. It opens the door for remote monitoring via SNMP (potentially from any place in the world) to improve the productivity of the network users.

From a cabling perspective, deployment of EtherNet/IP is similar to the deployment of DeviceNet; validation of proper connection of the physical layer is vital. Major differences, however, between EtherNet/IP and DeviceNet include the following:

- EtherNet/IP has a star topology; DeviceNet uses a trunk-line/drop-line topology.
- The switching infrastructure of EtherNet/IP becomes a part of the network that may need to be configured to optimize its performance. DeviceNet can be configured as a master-slave or as a distributed control architecture using peer-to-peer communication.

5.1 Validating the Integrity of the Physical Layer

Implementing the physical layer involves four steps: design, installation, cable verification, and certification or acceptance testing. Performance testing and verification of the EtherNet/IP physical layer infrastructure, whether copper or fiber-optic cable, is essential after installation and before network startup. This is similar in practice to testing any industrial network infrastructure, such as DeviceNet or ControlNet.

Each cabling segment (consisting of cable and connectors) must be tested to confirm that, after installation, the segments all conform to *The EtherNet/IP™ Specification*[21] for performance. As noted earlier, the use of pre-terminated cables is strongly advised to limit problems with field terminations. If commercially-supplied cables with RJ-45 connectors must be used, they should be kept in tightly-closed

[21] ODVA, The CIP Networks Library, Vol. 1 and 2, *The EtherNet/IP™ Specification.* Ann Arbor: ODVA, Inc., 2006. www.odva.org. CD-ROM.

cabinets to protect them from shop-floor environments. For critical applications, the use of more durable sealed connectors will stand up to the rigors of plant-floor environments.

Testing is easily done with commercially available hand-held network testers. Special adapters may be necessary for sealed connectors. Testing output includes conformance to all electrical requirements including, but not limited to, attenuation, impedance, return loss, cross-talk, and cable segment length measurements.

5.1.1 Link and Other Diagnostic LEDs

When deploying EtherNet/IP networks or diagnosing existing networks, engineers can take advantage of a few low-level, built-in diagnostic “tools” found on switches, routers, hubs, and network adapter cards on computers or industrial equipment. One of these tools is the built-in link/activity LED.

The link LED indicates a connection between two EtherNet/IP devices. These can be two switches, a switch and a computer, a switch and a PLC, etc. How the device knows there is a proper link is that each EtherNet/IP device, when not transmitting data, will transmit a link signal on all of its EtherNet/IP ports. When either device receives this link signal, it will then turn on the link LED associated with that port. Therefore, if the wiring is correct between two EtherNet/IP device ports, the appropriate link LEDs will be enabled on each of the devices at both ends of the cable.

There may also be speed LEDs and type of duplex connection LEDs that indicate the speed (10 Mbps, 100 Mbps, etc.) and the type of duplex connection (half- or full-duplex) between the two devices. Some devices may use a multicolor LED to indicate speed. For example, amber may indicate 10 Mbps with green for 100 Mbps. Both sides of the connection must use the same speed and the same duplex settings for proper communications. Most equipment will automatically negotiate for the fastest speed and for full-duplex operation. If both devices connected to a cable show a 10-Mbps and half-duplex setting when they are rated for 100-Mbps and full-duplex operation, this may be an indication of a defective cable. It is also possible that there may be a physical problem with the adapter’s circuit at either end, or a configuration switch may be set for manual operation, locking in the speed or mode of operation.

Many switches and EtherNet/IP devices also provide an activity LED to indicate data transmission activity. Sometimes the link LED flashes to indicate either receive or transmit activity. The manufacturer’s manual should be reviewed for specific details.

The link LED will indicate a working connection between two EtherNet/IP devices. The activity LEDs will show that one of the two devices is transmitting data. The speed and duplex LEDs can provide further diagnostic information. Sometimes a single LED will be used to show multiple functions. In some cases a router or switch may even have separate LEDs to show transmit and receive (or upload and download) functions. In many cases, there is not consistency from device to device. It’s best to check the manual to confirm what the various LEDs indicate.

In rare cases, link LEDs can show a proper connection between two devices even though the cable may be constructed improperly. Communication may still occur over an improperly wired cable, but it may be intermittent or contain errors. A cable tester tool should be used to check all cables to evaluate fully the cable wiring. Fiber-optic cables can have a similar problem, especially if they have to be terminated in the field under poor lighting or temperature conditions.

5.1.2 Ping Checks for TCP/IP Continuity

If proper wiring between two devices has been verified, another built-in diagnostic tool that can be used is the ping. Like the link LEDs built into EtherNet/IP devices, ping is built into TCP/IP software.

A ping is basically an Internet Control Message Protocol (ICMP) echo request and echo response between two TCP/IP devices. As EtherNet/IP uses TCP/IP, EtherNet/IP devices may support ping. Managed switches may also have the ability to respond to ping and may be able to send a ping request.

The device sending the ICMP echo request (also called a ping request) will transmit the request, and the addressed device will provide the echo response or ping response. Ping requires that both devices have proper IP addresses and that both exist on the same subnet or are connected via one or more routers.

Most all computers—whether Linux/UNIX, DOS, or Windows—support ping, usually from the command line, and sometimes from a Graphical User Interface (GUI) application that supports pings. For example, a Windows 95 or later computer with a proper network configuration can both transmit a ping request and respond to a ping request. In an MS-DOS or command-line window, the user types [ping] followed by the IP address of the targeted device as shown in Figure 5-1. In Windows, the ping application will transmit four ping requests; additional or fewer pings can be requested using [ping –n x 192.168.25.25] where x is the number of pings to be sent to the address that follows. The resulting output will show the time it took for each response, or it will indicate that the request timed out if the ping was unsuccessful.

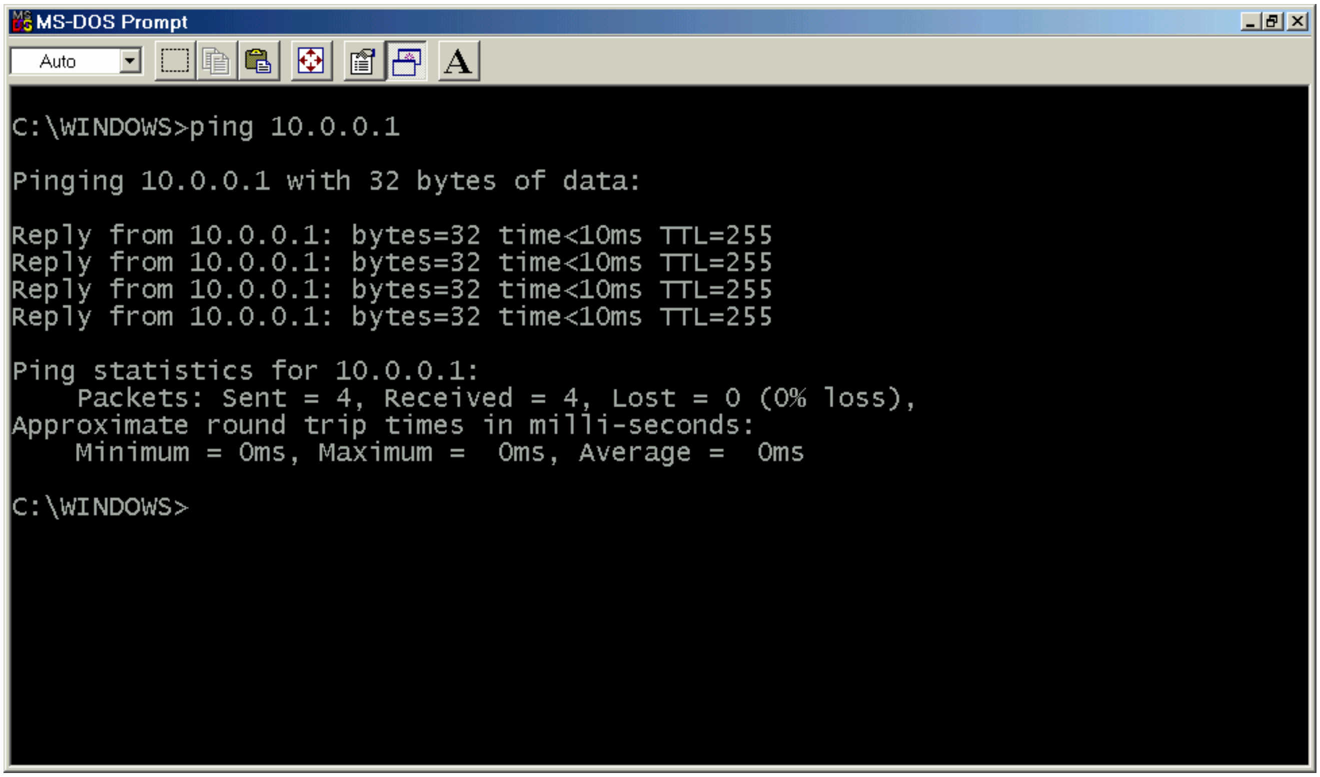

```
C:\WINDOWS>ping 10.0.0.1

Pinging 10.0.0.1 with 32 bytes of data:

Reply from 10.0.0.1: bytes=32 time<10ms TTL=255
Reply from 10.0.0.1: bytes=32 time<10ms TTL=255
Reply from 10.0.0.1: bytes=32 time<10ms TTL=255
Reply from 10.0.0.1: bytes=32 time<10ms TTL=255

Ping statistics for 10.0.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum =  0ms, Average =  0ms

C:\WINDOWS>
```

Figure 5-1 Simple Ping Request Sent from the Windows Command Line Using an MS-DOS Prompt.

Linux and UNIX computers can send pings from the command line as well, but if a count is not specified, the computer will continue to send pings until a terminate command is given by hitting [Ctrl-c]. To specify a number of pings under UNIX or Linux, type [ping –c x 192.168.25.25] where *x* is the number of pings to be sent to the address that follows.

Ping can be used to validate connections and IP addresses and to help determine where connections are failing. To help locate faults, an engineer can ping end devices, devices located on intermediate points, or managed switches between two end devices.

Most TCP/IP devices will respond to a ping request, however, it is possible for a device to ignore ping requests or to have ping responses disabled. The device's user manual will usually note if pings or ping requests can be turned off.

5.1.3 Switch Configuration and Management

EtherNet/IP networks can be optimized to enhance their performance. Most industrial EtherNet/IP deployments that share the network with other kinds of traffic will require some level of switch configuration. The configuration process also enables the network to insure a high level of resiliency, deterministic behavior, and a high degree of security with appropriate access for authorized users. This process will increase the level of access an authorized user has, potentially allowing for monitoring and configuration of devices from anywhere in the world. It also reduces costs as it leverages the existing Ethernet infrastructure to deliver multiple services and new applications.

Managed vs. Unmanaged Devices

The broad definition of a managed device encompasses any network element that can be configured to improve performance and/or can be accessed remotely to gather metrics on it. Switches that can be configured and accessed remotely are normally referred to as managed switches. Network switches that are not configurable, and therefore cannot be optimized, are generally referred to as unmanaged switches.

In terms of the EtherNet/IP infrastructure, managed devices may include very basic functions such as simple network utilization (RMON) to more critical settings that determine who has access to a given set of services or traffic (security). It can also include which traffic should be prioritized (QoS), what amount of bandwidth should be allocated to a given port or application to assure high availability of the network (preventing potential critical situations like a broadcast storm from interfering with control messaging), how to deal with multicast traffic, etc. The level of management is commensurate with the set of features that may be optimized or the metrics that may be gathered.

Many network elements, such as routers and switches that have been designed as managed devices, are configured using Simple Network Management Protocol (SNMP), often with graphical user interfaces. This is actually no different from programming a PLC or any other control element that has some level of embedded programmable logic.

Most vendors of Ethernet switches provide some graphical interface utility that can be used for accessing the device. Many of the vendors also provide a web browsing interface allowing for the configuration of the devices through the network by pointing the browser to the switch's IP address. This presumes that the IP address has already been configured into the device.

The IP address of the switch can generally be configured using a console cable directly connected to the device. Alternatively, a Dynamic Host Configuration Protocol (DHCP) or Bootstrap (BootP) server can be used.

Managed switches also allow for remote management of the network. There are multiple network management software vendors that provide solutions to manage Faults (alarms) remotely, Configuration, Accounting (who is using what resources), Performance, and Security—commonly referred to as FCAPS.

Problems identified during infrastructure testing will be detailed by the network tester so appropriate corrective action can be taken. Other non-conformance problems and warnings may appear that address additional issues, such as incorrect termination, faulty components, and high ambient noise. All non-conformance and test warning issues must be resolved before network startup.

5.2 Putting Traffic on the Network

In order to identify features of the EtherNet/IP network infrastructure that help provide the required performance and connectivity, it is necessary to characterize network traffic, employ IP addressing techniques, and understand IP multicast traffic.

5.2.1 Characterization of EtherNet/IP Traffic

As has been shown, EtherNet/IP network infrastructure can be defined as a hierarchical interconnection of Layer 2 and Layer 3 Ethernet switches.

Traffic generated during programming, configuration, and diagnostics of EtherNet/IP devices as well as during exchange of non-time-critical data is called explicit messaging in EtherNet/IP terminology. It is normally low-rate traffic that generally has an insignificant impact on network performance. Although it contains both broadcast and unicast data, this traffic does not require engagement of any special features in the EtherNet/IP network infrastructure. Broadcast traffic typically consists of IP packets supporting Address Resolution Protocol (ARP), BootP, DHCP, SNMP and other protocols of this type. Unicast traffic consists of Transmission Control Protocol/Internet Protocol (TCP/IP) packets.

Traffic generated during time-critical data exchange is called implicit messaging and consists mostly of User Datagram Protocol/Internet Protocol (UDP/IP) unicast and multicast packets. Examples include:

- Input/Output (I/O) data and status produced by a remote I/O device for consumption by one or more programmable controllers
- Data produced by a programmable controller for consumption by one or more programmable controllers

While the handling of UDP/IP unicast traffic does not require engagement of any special features in the EtherNet/IP network infrastructure, handling of UDP/IP or IP multicast traffic may require these features. Although EtherNet/IP supports change-of-state and application reporting, in a typical control system, data exchange is predominately cyclic. The implicit messaging traffic is normally generated at an aggregate rate of tens of thousands of packets per second, depending on the number and type of EtherNet/IP devices and the application. Many EtherNet/IP devices are, for example, capable of generating up to 5000 packets per second. Normally, this traffic is evenly divided between UDP/IP unicast and multicast packets. Packet length is typically around 120 bytes.

Normal IP multicast traffic generated in an EtherNet/IP network consists of high-rate, short-packets generated on a continuous basis. For this reason, EtherNet/IP networks differ considerably from typical office networks, where IP multicast traffic is generated sporadically and with much lower packet rates. A growing exception to this traffic profile may be in the area of multimedia audio and video conferencing applications.

EtherNet/IP Specification Tip

In accordance with *The EtherNet/IP™ Specification*, a message producer can send data encapsulated in either an IP unicast or IP multicast packet.

5.2.2 IP Addressing in EtherNet/IP Networks

In some EtherNet/IP products, transfer of CIP-implicit messages is performed using connectionless transport layer protocol, UDP, and multicast services of underlying protocols, namely IP and Ethernet. Only "heartbeats" are transferred as UDP unicast packets. Therefore, the first half of a transaction constituting implicit message transfer from a producer to a consumer is a UDP/IP multicast packet transmitted by the producer, and the second half is a UDP/IP unicast packet transmitted by the consumer. In the case of multiple consumers, each consumer transmits a "heartbeat" packet to the producer.

If an EtherNet/IP bridge or adapter has a maximum packet rate of 5000 packets per second, it means that a maximum of 2500 packets per second can be transmitted by this device, and a maximum of 2500 packets per second can be received by it. Therefore, the maximum UDP/IP multicast packet rate of this device is 2500 packets per second.

It is important to remember that the destination address of an IP multicast packet is always selected from a specific range, called Class D and, in the case of some EtherNet/IP products, from a specific sub-class within it. Therefore, as an example, these products generate IP multicast addresses from 239.192.xxx.xxx to 239.193.xxx.xxx.

5.2.2.1 IP Addressing Example

Consider an EtherNet/IP network consisting of a programmable controller system acting as a consumer, a remote I/O rail acting as a producer, and a managed switch. Assume that the consumer has been configured with the IP address 10.88.80.100, the producer has been configured with the IP address 10.88.80.110, and the switch has been configured with the IP address 10.88.80.120. The CIP connection has been configured with a 10-millisecond (ms) Requested Packet Interval (RPI). See *6.1.1 Multicast Traffic Management within a Control System* for more information.

During the CIP connection establishment process, the consumer sends a CIP Forward_Open Request packet to the producer. This is a TCP/IP unicast packet with IP Destination Address 10.88.80.110 and IP Source Address 10.88.80.100. In response, the producer sends a CIP Forward_Open Response packet back to the consumer, which contains an IP multicast address in its Data field that the producer will use to send data to the consumer on this CIP connection. The Forward_Open Response packet is also a TCP/IP unicast packet but with IP Destination Address 10.88.80.100 and IP Source Address 10.88.80.110.

Assume that the IP multicast address generated by the producer is 239.192.1.1. As soon as a CIP connection has been established, the producer begins multicasting a UDP/IP packet containing the status of its inputs every 10 ms. The IP Destination Address of this packet is 239.192.1.1, and the IP Source Address is 10.88.80.110. The consumer responds by sending a "heartbeat" packet, which in this example contains the output status update. This packet is a UDP/IP unicast packet with the IP Destination Address

10.88.80.110 and the IP Source Address 10.88.80.100. The packet rate will be 200 packets per second, consisting of 100 multicast and 100 unicast packets per second.

5.2.3 IP Multicast Traffic in EtherNet/IP Networks

To satisfy the performance requirements of the implicit messaging traffic, the following is recommended:

- EtherNet/IP network infrastructure should be based on Ethernet switching technology.
- Micro-segmentation should be used.
- Networks should run in 100-Mbps, full-duplex mode.

Ethernet switches (also called Layer 2 or Ethernet Layer 2 switches) should always be used instead of hubs. Only one device (end-node) should be connected to a switch port, and, if manual-mode setup is preferred, switch and device ports should be set to 100-Mbps, full-duplex mode. Switches should support port mirroring and IGMP snooping. Port mirroring is convenient for diagnostic and troubleshooting and is covered in *Appendix A Recommendations for Ethernet Switches in EtherNet/IP Systems*. IGMP snooping helps to manage multicast traffic and is described in *6 Infrastructure Application Scenarios* in more detail. Without the use of special traffic filtering features, an Ethernet Layer 2 switch normally retransmits each received IP multicast, broadcast, or unknown unicast packet through all ports.

5.3 Deployment Summary

In many ways wiring, testing, and troubleshooting an EtherNet/IP network is not unlike wiring a complex telephone system. Cable lengths are more important, and proper termination is a must. In addition to simple hand-held testers, most of the infrastructure components—routers, switches, and other equipment—provide built-in testing capabilities for both the physical wiring and for TCP/IP testing. Checking for obvious potential problems before putting the system into use is important. These problems include cable routing mistakes, environmental concerns, and location of equipment. Once the wiring is functioning, the guidelines presented here can be used to check and maintain traffic performance and flow.

6 Infrastructure Application Scenarios

Regardless of the specific application, all networks require a design that deals with multicast traffic to meet the needs of real-time control.

EtherNet/IP supports both the time-critical (implicit) and non-time-critical (explicit) message transfer services of CIP. With CIP, the exchange of time-critical messages (real-time control or I/O control) is based on the producer-consumer model where a transmitting device produces data on the network, and many receiving devices consume this data simultaneously. On EtherNet/IP, this functionality is provided through IP multicast and Ethernet multicast mechanisms. The use of Ethernet for real-time control is different from typical office and other non-time-critical applications.

In the office world, one of the primary network design objectives is to manage the traffic load on the network. In the industrial world, the small, embedded control devices have modest CPU capability compared to the office world's PCs and workstations. Device CPU resources are needed to process each broadcast and multicast message. While managing network bandwidth is always important, correct design and setup of the Ethernet infrastructure must first protect the control devices from multicast and broadcast floods that can slow or stop operation of these devices.

The recommendations in this section focus on addressing real-time control application needs. These applications typically use the producer-consumer capabilities of EtherNet/IP, which are based on multicast Ethernet communications. Monitoring, device configuration management, or explicit message applications that use unicast communications do not need to follow the recommendations for multicast traffic outlined in the following scenarios.

There is no single or standard approach to infrastructure implementation will provide the right balance of performance, data connectivity, and installed cost for all applications. However, these scenarios address the general application concerns of four typical types of industrial EtherNet/IP networks.

- Isolated EtherNet/IP network with a single controller. An isolated EtherNet/IP network interconnecting devices of a control system with a single controller.

- Isolated EtherNet/IP network with multiple controllers. An isolated EtherNet/IP network interconnecting devices of a control system with multiple controllers.

- EtherNet/IP network connected to the enterprise network. An EtherNet/IP network connected, via a gateway or equivalent, to the enterprise network. The control system it serves can be of any complexity.

- EtherNet/IP network integrated with the enterprise network. An EtherNet/IP network that is fully integrated with the enterprise network. The control system it serves can be of any complexity.

Within each of these four network types, recommendations can be made regarding the types of active infrastructure components (switches, routers, etc.) and the key functions needed to balance real-time performance with data connectivity requirements. *Appendix A Recommendations for Ethernet Switches in EtherNet/IP Systems* complements this chapter by summarizing the recommended switch functions for each of the four system types.

6.1 General Application Considerations

There are several key concerns that vary based on the type of system desired.

6.1.1 Multicast Traffic Management within a Control System

Data producers on an EtherNet/IP network can generate a significant amount of multicast traffic, requiring that particular care be taken to manage it. Managing multicast traffic is the most important concern in the design of a high-performance EtherNet/IP network. With EtherNet/IP, implicit messages (I/O control and controller-to-controller communication) use unicast (one-to-one) packets to communicate output information from a controller to I/O or other control devices. Input information, sent from these devices to the controller, is usually transferred using multicast (one-to-many) packets. As noted in *4.7 Switching Technologies*, switches normally retransmit multicast packets to all switch ports (flooding), so that all devices see all the multicast traffic. Without multicast control functions, such as IGMP snooping, the end devices' operation is impaired as its CPU becomes overloaded due to discarding extraneous multicast traffic.

Requested Packet Interval

One of the major controller or scanner setup parameters is the Requested Packet Interval (RPI). This value is the minimum rate at which control devices will produce the input information needed by the controller. The larger the application (and the number of interconnected control devices) and the faster the response time needed (lower RPI), the greater the need for IGMP snooping.

The tradeoff here is that managed switches (those typically with web/SNMP diagnostics and message filtering functions, such as IGMP snooping, VLAN, and others) cost more than simple unmanaged switches, which do not have message filtering functions, such as IGMP snooping. Systems with the highest performance requirements and the fewest application timing problems will use all managed switches with IGMP snooping.

6.1.2 Control Device Sensitivity to Extraneous Traffic

Another consideration related to multicast message control relates to each control device's sensitivity to unwanted multicast and broadcast traffic. The impact of unwanted traffic is analogous to sorting the "junk mail" from the important mail.

6.1.2.1 Potential for Traffic Overload

Every time a control device receives an Ethernet frame, it takes a tiny bit of its microprocessor/Ethernet chip resource to look at the address, decide if the message is for it, and discard messages addressed to others. The rate of incoming unwanted messages (expressed in packets per second) can reach a point where so much time is spent processing and throwing out junk mail that it can not produce or receive control messages at the configured RPI value. This will cause an error condition that shuts down communications and causes the system to go into its fault mode.

6.1.2.2 Varying Sensitivity by Device Type

Certain devices, such as Windows-based PCs/workstations, may have an extra sensitivity to multicast traffic. There can be significant variance in sensitivity between the same devices from different vendors, and between different devices from the same vendor. Direct connection of the sensitive devices to separate Ethernet ports on controllers or to Ethernet connections via switches capable of IGMP snooping may be mandatory.

6.1.3 Configurable Isolation of Enterprise and Control Traffic

One of the promises of Ethernet-based systems is *shop-floor-to-top-floor connectivity*. When connecting control networks to enterprise networks (i.e., networks supporting mainly office traffic, MES, plant management, manufacturing line efficiency, quality monitoring , etc.), unwanted enterprise traffic (including multicast and broadcast) must be blocked from disrupting control traffic. Equally important is the need to prevent EtherNet/IP multicast control traffic from flooding and bogging down the enterprise network. The conventional method used to connect separate fieldbus networks is to use two network interface cards in a PC. The enterprise and control networks are physically isolated with a PC-based program used to selectively enable data transfers in both directions.

Of course, any change in the desired data flow requires each of these PCs with two network cards to be manually reprogrammed. For single connections and relatively stable data flows, this may be a completely acceptable approach. However, as the quantity of control-to-enterprise connection points grow and the need to adapt data flow increases, reprogramming all the intermediate databases becomes unwieldy. In these cases, industrial switches and routers with Virtual LAN (VLAN) capabilities allow configurable traffic isolation without reprogramming PCs, rewiring network cables, or purchasing additional network cards. VLAN-capable switches and routers isolate broadcast and multicast traffic, preventing it from passing into other network areas. These same functions and devices are also used to implement network security functions. See *4.10 Virtual Local Area Networks* for more VLAN information.

Once traffic is isolated using VLANs, routers (or Layer 3 switches with routing functions) are required to allow data to be selectively transported from one VLAN to another. Other router mechanisms, such as Time-to-Live (TTL) thresholds and Access Control List (ACL) filtering may be used in conjunction with VLANs to provide further isolation and protection between plant-floor and enterprise environments.

6.2 Isolated Control Network with a Single Controller

Isolated EtherNet/IP networks are not connected to other plant or enterprise-level networks. Control system examples built around these networks typically contain one programmable controller, one or more workstations, and control devices, which may include I/O nodes and/or drives, etc. Two configurations are typical: low (usually with 10 or fewer devices) and high (usually with more than 10 devices). The major performance concerns relate to managing multicast traffic and preventing device overload from unwanted multicast traffic, which would shut down device communications.

6.2.1 Systems with Low Device Counts

In the example shown in Figure 6-1, IP multicast traffic is produced by I/O devices and consumed by the programmable controller. Here, IP multicast traffic produced by I/O_1 will be sent to all devices connected to the switch, including I/O_2 through I/O_5 and the workstation.

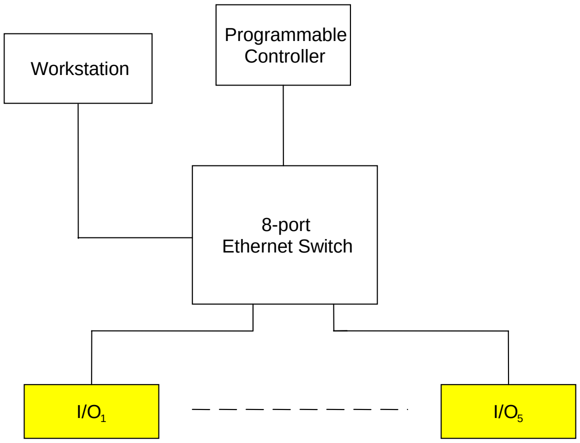

Figure 6-1 Isolated Network with One Controller.

EXAMPLE 1

Assume that each device establishes one CIP connection to the controller with a 10-ms RPI. Each device will generate 100 multicast packets per second and receive 100 unicast ("heartbeats" or output status) packets per second from the controller. Each device will also receive 400 unwanted multicast packets per second, and the workstation will receive 500 unwanted multicast packets per second.

EXAMPLE 2

Assume that each control device establishes four CIP connections with the controller. One connection has a 10-ms RPI, and each of the other three connections has a 50-ms RPI. Each control device will generate 160 multicast packets per second and will also receive 640 unwanted multicast packets per second. The workstation will receive 800 unwanted multicast packets per second.

6.2.1.1 Use of Unmanaged Switches

In isolated systems with low device counts (10 or fewer), a low number of CIP connections, and modest RPIs (higher than 10 ms), the level of unwanted multicast traffic may not cause a device overload condition. In isolated systems with large device counts and, more importantly, a large number of CIP connections, the use of unmanaged switches as expansions to managed switches may allow a balance of cost versus performance. In applications with a low number of CIP connections that have modest RPI requirements, the use of unmanaged switches can provide product cost savings over the use of all managed switches. As a general rule, applications requiring fast RPIs and/or more than a few connections should use managed switches with IGMP snooping capability.

When using unmanaged switches, follow these guidelines:

Calculate or test the maximum multicast traffic load for each device.
For each device, ask the vendor for the maximum unwanted multicast packets per second rate that the device can withstand while maintaining the desired RPI rate. Be sure to use the configuration with the largest I/O count or most time-critical configuration. If the vendor cannot provide this information, refer to the alternative traffic-management approaches presented below.

For each device connected by one or more unmanaged (non IGMP snooping) switches, add all the rates together to determine the total multicast traffic in the multicast domain, then add in any significant broadcast traffic from application or network management software, and reduce the total by 5% for a design buffer. Note that splitting the devices between different unmanaged switches makes no difference in the total multicast domain traffic load as depicted in Figure 6-2.

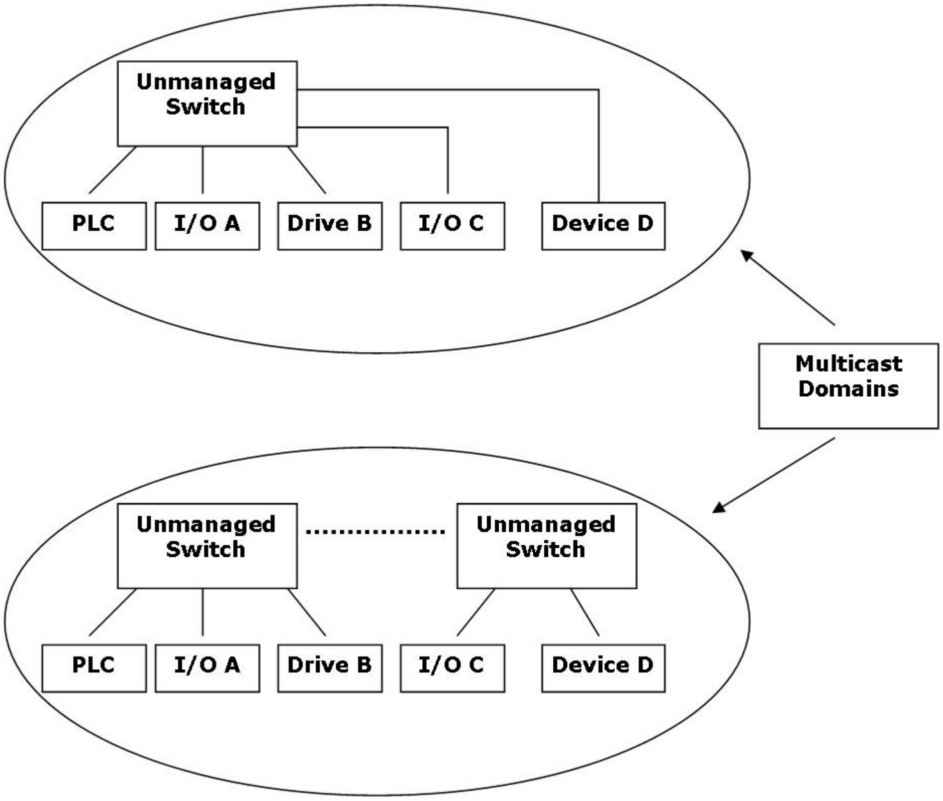

Figure 6-2 Adding Unmanaged Switches Does Not Reduce the Size of a Multicast Domain.

Confirm that the total unwanted multicast traffic value is less than each devices' maximum unwanted multicast traffic rate. If it is greater, add managed switches. Adding managed switches reduces the multicast domain as shown in Figure 6-3.

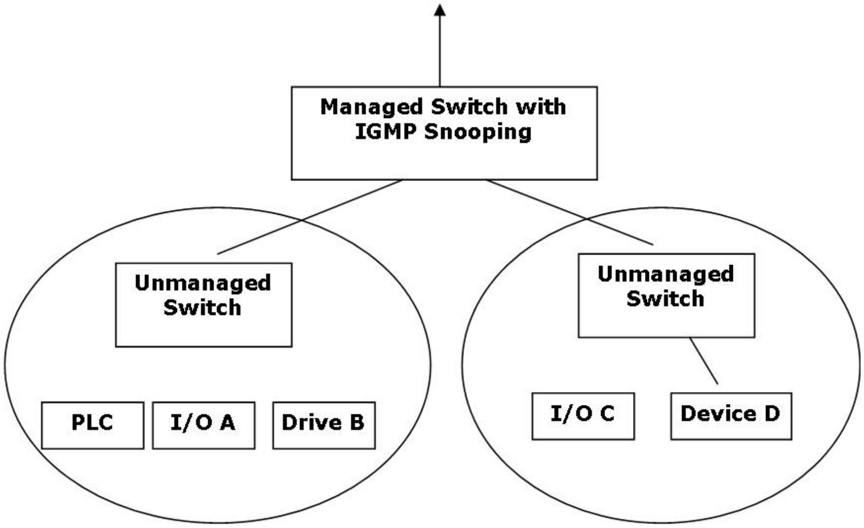

Figure 6-3 Use Managed Switches to Reduce the Size of a Multicast Domain.

This is also a key concept for end users who may be interconnecting multiple OEM machines from one or more vendors. To minimize application timing problems and maximize determinism, separate the traffic from the different machines with managed switches.

Consider future expansion needs.

In the future, before connecting any additional devices to the unmanaged switches, first determine if the unwanted muticast traffic load of the existing installation (as well as the added device) can withstand the added device's traffic. Failure to do this runs the risk of adding one device that shuts down the system. It is reccommended that any new sections or expansions be connected into a managed switch to add IGMP snooping capabilities.

Alternative traffic management approaches.

In cases when the prefered calculation approach is not possible and the maximum unwanted multicast traffic rates cannont be obtained from the device vendors, there are three options:

1. Use all managed switches (with IGMP snooping). This is the best alternative and the most conservative approach to system design because it keeps traffic and application timing issues to a minimum.

2. Test the configuration during the design phase. Pre-test each multicast domain (unmanged switch section) to confirm proper operation using the actual application configuration and controller/scanner RPI value. Be sure to allow time to add managed switches to the system design if they are needed.

3. Reconfigure devices to use unicast messages. Some devices may have a user-selectable option to produce implicit (I/O) messages using unicast messaging instead of

multicast messaging. Disabling multicast messages allows the use of unmanaged switches. This is most effective for systems with one consumer.

In the systems described in these examples, a certain amount of the workstation's and remote I/O devices' computing power is used to filter the unwanted traffic. If the control engineer concludes that unwanted multicast traffic affects the performance of these devices and, consequently system performance, then one of the following recommendations may be considered:

Certain types of devices, such as PCs/workstations, can be more sensitive to unwanted multicast traffic. If only workstation performance is affected, the controller should be connected via a separate switch as shown in Figure 6-4. IGMP snooping need not be activated in any of these switches.

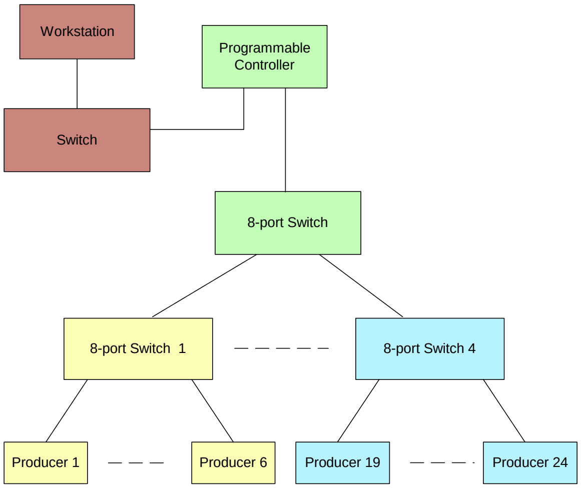

Figure 6-4 Isolated Network with One Controller and a Separate Network for HMI.

If both workstation and remote I/O performance is affected, a switch with IGMP snooping should be used. The switch with IGMP snooping forwards IP multicast packets to consumer ports only. In isolated multi-switch networks, one switch with the IGMP query function is needed, and the other switches need to support the IGMP snooping function.

6.2.2 Systems with High I/O Count

Isolated systems with high producer device counts (i.e., I/O, drives, sensors, etc.), a large number of CIP connections, and low RPI rates require the use of managed switches for implicit-messaging-based I/O control applications. Larger non-time-critical applications (such as monitoring systems) with RPI rates over 70 ms may be capable of using combinations of managed and unmanaged switches. Use the guidelines in *6.2.1.1 Use of Unmanaged Switches* to determine the potential for multicast overloads. As systems grow larger, the need for web or SNMP diagnostics can considerably reduce initial system startup times, and provide port-based traffic statistics to help optimize overall system performance. In addition, as systems grow in complexity and size, the need for network redundancy often grows as well.

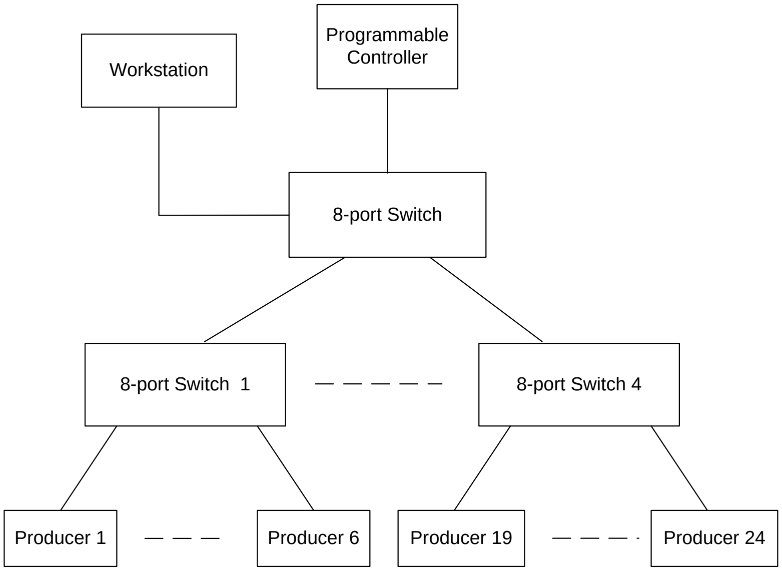

Figure 6-5 Isolated Network with One Controller.

EXAMPLE 3

In the example shown in Figure 6-5, IP multicast traffic is produced by 24 producers, remote I/O rails and/or drives, and is consumed by the programmable controller. Assume that each producer establishes one connection with the controller with a 10 ms RPI, and IGMP snooping is not activated in the switches. Each producer will generate 100 multicast packets per second and will receive 2300 unwanted multicast packets per second. The workstation will receive 2400 of unwanted multicast packets per second.

If the engineer concludes that unwanted multicast traffic affects system performance, then one of the following recommendations may be considered:

- If only workstation performance is affected, the workstation can be connected to the controller via a separate switch as shown in Figure 6-4. IGMP snooping does not have to be activated in any of these switches.

- If workstation and remote I/O performance are affected, a switch with IGMP snooping should be used. In isolated networks, one of the switches will be required to have the IGMP query function to manage the multicast process as shown in Figure 6-6.

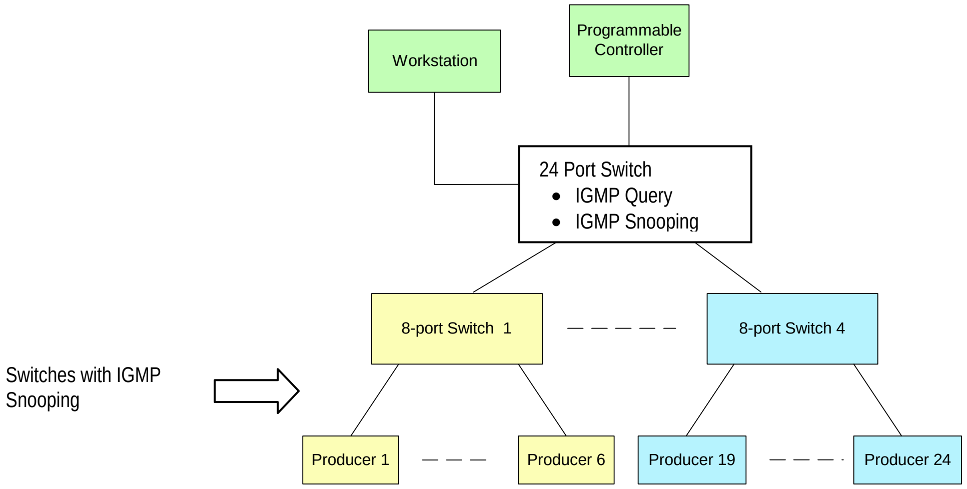

Figure 6-6 Isolated Network with One Controller and a Main Switch with Query Function.

EXAMPLE 4

Figure 6-7 depicts a "daisy chain" configuration similar to and operating under the same conditions as the one shown in Figure 6-5.

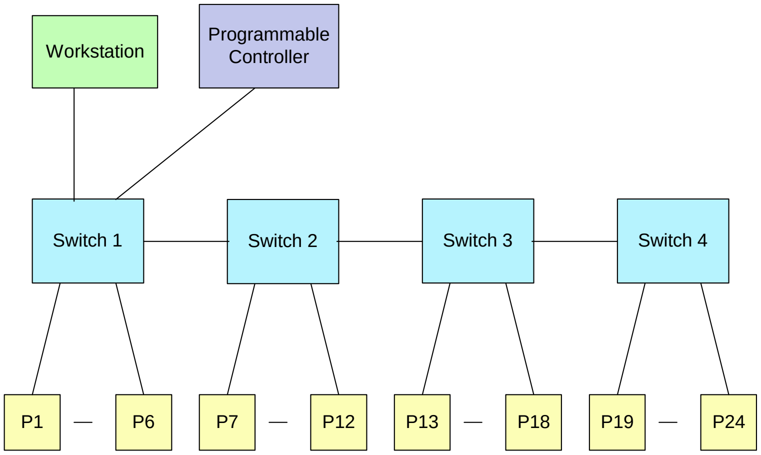

Figure 6-7 Isolated Network with One Controller and a Daisy Chain Configuration.

Here, each producer establishes one connection with the controller with 10-ms RPIs. Each producer will generate 100 multicast packets per second and will receive 2300 unwanted multicast packets per second. The workstation will receive 2400 unwanted multicast packets per second and each switch will process 2400 multicast packets per second.

In this example, if producer traffic is `unicast` instead of multicast, then `switch 4` will process 600 packets per second, `switch 3` will process 1200 packets per second, `switch 2` will process 1800 packets per second, and `switch 1` will process 2400 unicast packets per second. It is important to note:

- In this unicast environment, no matter how many packets a switch will process, no nodes will receive extraneous traffic, thus eliminating the potential of a multicast traffic loading problem.
- In a daisy-chain connection of switches in this unicast environment, switch 1 has very high loading. To reduce the additive loading from switch to switch, a star topology with switch 1 in the middle is recommended instead of the daisy chain configuration. This will also decrease the individual port loading on switch 1.

In this example, if the produced traffic is `multicast`, and both workstation and remote I/O performance are affected, switches with IGMP snooping should be used. To even loading across the network, a star topology is used instead of a daisy-chain topology.

6.3 Isolated Network with Multiple Controllers

EtherNet/IP networks can consist of multiple programmable controllers (or other controllers such as motion) and still remain isolated—not connected to other networks. One example is the VLAN approach shown in Figure 6-8 and Figure 6-9 where two programmable controllers use a single switch to connect to their I/O.

6.3.1 Multiple Single-Controller Systems without Data Sharing

Figure 6-8 depicts two small, separate control systems sharing the same managed switch. Programmable Controllers 1 and 2 are not sharing data. Also, each of the I/O devices produces data for only one programmable controller. In other words, this represents two single-consumer systems that share a managed switch. To prevent one system’s unwanted multicast traffic from appearing on the other system, managed switches with VLAN functionality must be used.

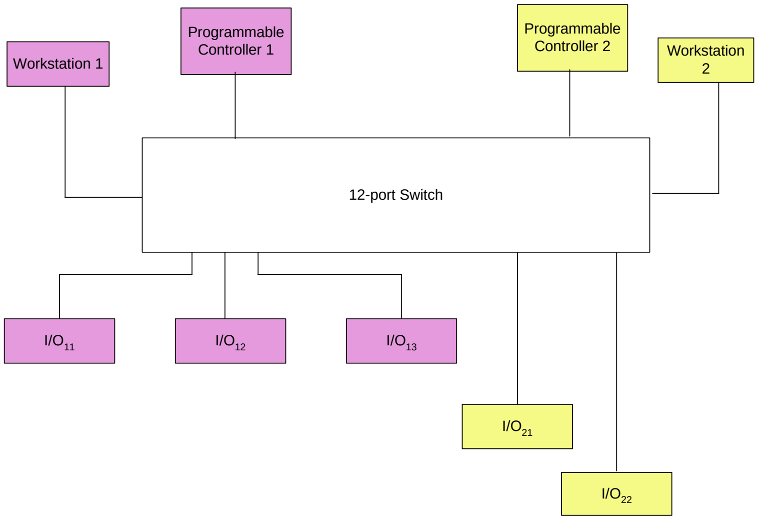

Figure 6-8 Two Single-Controller Systems Sharing a VLAN-enabled Managed Switch.

Figure 6-9 depicts the logical view of the system set up as two VLANs. Here, ports 1, 3, 4, 5, and 6 belong to VLAN 1; ports 2, 7, 8 and 9 belong to VLAN 2. For more distributed applications, VLANs can also be implemented across multiple switches.

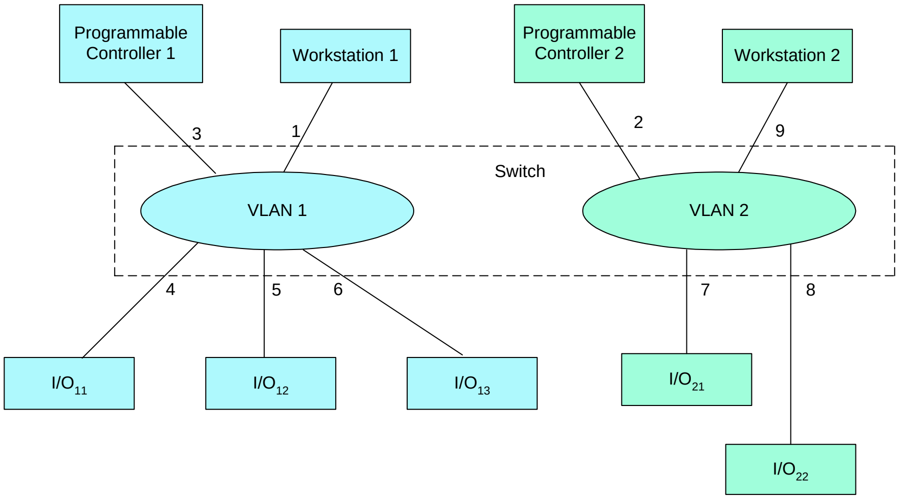

Figure 6-9 Logical View of Two Single-Controller Systems Sharing a VLAN-enabled Managed Switch.

6.3.2 Multiple Single-Controller Systems with Data Sharing

Consider the example shown in Figure 6-9 but with the need for data sharing. If programmable controllers need to share data or if input data produced by one more of the I/O devices needs to be shared, VLANs are not an option. Infrastructure requirements for a network serving these applications are the same as for an isolated network with a single controller as described in *6.2.2 Systems with High I/O Count*. Even with one switch, the number of devices connected to the network and the number of CIP connections demand IGMP snooping and query to prevent a multicast overloading of devices.

6.3.3 Isolated Systems Summary

In small isolated EtherNet/IP networks with a single controller, control of IP multicast traffic using IGMP snooping may not be required. In these cases, unmanaged switches may be used, given that the overall multicast traffic load is calculated or the system pre-tested to prevent device shutdown from multicast traffic overloads. In cases when unwanted multicast traffic affects performance of HMI workstations, a separate network for controller-HMI communication should be used.

In large isolated networks when unwanted multicast traffic also affects the performance of traffic producers (remote I/O, drives, etc.), managed switches with IGMP snooping could be used in combination with unmanaged switches to lower total system cost. For maximum system performance, all managed switches with IGMP snooping should be used. When IGMP snooping is used, one of the switches must support the IGMP query function.

In isolated networks with multiple controllers that do not share data between them, managed switches with VLAN functionality may be used to increase control performance by isolating traffic. If data sharing between controllers (i.e., implicit messaging) or the use of produce-consume tags is required, infrastructure requirements are the same as those for an isolated network with a single controller. The same is true when controllers share remote inputs.

6.4 Enterprise-Connected and Integrated Control Systems

There are often requirements for enterprise or plant-wide networks to selectively access control system data. Higher level plant systems often contain servers used to download the control programs necessary to reconfigure flexible manufacturing systems. Plant-wide systems may also carry video monitoring, MRP, and other production data that needs to be processed separately from real-time critical control traffic on the plant floor. There are two approaches to allowing this selected access to occur while isolating the plant-floor network from the plant backbone.

6.4.1 Networks Connected via CIP Gateways

An EtherNet/IP network can be connected to the enterprise network via a gateway. A typical example of a gateway is a programmable logic controller with multiple EtherNet/IP ports acting as a CIP gateway as shown in Figure 6-10. In such a network, all of the implicit real-time control traffic is isolated to one or more separate EtherNet/IP ports. Since EtherNet/IP and enterprise networks are connected via a CIP gateway, there is no propagation of unwanted traffic from the EtherNet/IP network into the enterprise network and vise versa. Therefore, there is no difference, from the EtherNet/IP infrastructure point of view, between enterprise-connected and isolated EtherNet/IP networks. This approach requires that any data transfer between networks be done through PLC programming or setup.

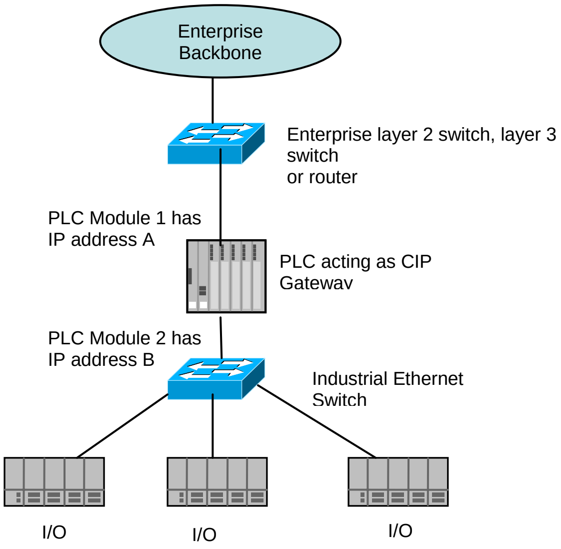

Figure 6-10 EtherNet/IP and Enterprise Networks Connected via a CIP Gateway.

6.4.2 Networks Integrated with Routers (or Layer 3 Switches)

An enterprise-integrated EtherNet/IP network is fully integrated into the enterprise network infrastructure. This type of interconnection results in a single, integrated network. The addressing practices, security policies, network management techniques, and “change control” procedures of the enterprise network and the control system portions of that network must be harmonized.

6.4.3 Two Types of Integration

The type of integration—non-interlock connected or interlock connected—used depends on whether or not implicit message traffic will traverse the corporate backbone network. These two integration types are representative of large corporate installations where the corporate backbone and the plant-floor environments span a large area. Non-interlock is quite common as it allows selected data transfers, but keeps the real-time data isolated to a particular area of the plant floor and isolated from other plant traffic that may affect control system performance.

6.4.3.1 Non-Interlock Connections

A non-interlock connection is used if there is no implicit message traffic traversing the backbone network. As shown in Figure 6-11, the only traffic being passed is supervisory, programming, monitoring, and so forth. There is no real-time data being passed. There is no CIP gateway in use; the backbone is connected directly to the plant-floor network infrastructure. There may be implicit messaging within the local network zone, but it is not meant to pass into the backbone network. Protection may be needed to prevent multicast traffic from adversely affecting the backbone environment. Connection points to the backbone may be a Layer 2 switch (VLANs, Layer 2 filtering), a Layer 3 switch (Access Control Lists, Layer 2/3 filtering), or a router (ACLs, Layer 2/3 filtering, firewall function).

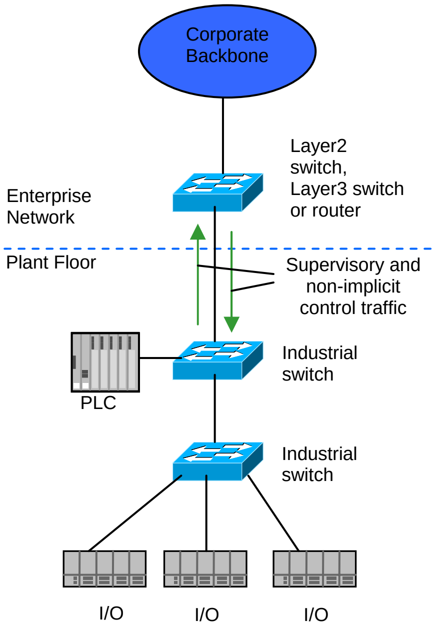

Figure 6-11 Example of a Non-interlocked Integrated Network.

6.4.3.2 Interlock Connections

An interlock connection is used if there is implicit message traffic traversing the backbone network. As shown in Figure 6-12, the traffic being passed is real-time implicit, supervisory, programming, monitoring, and so forth. There is real-time data being passed. There is no CIP gateway in use. The backbone is connected directly to the plant-floor network infrastructure. There will be implicit messaging within the local network zone and some will be passed across the backbone network to another zone. Protection (firewall, ACL, VLAN, etc.) may be needed to prevent unwanted multicast traffic from adversely affecting the backbone environment or the plant floor. The recommended connection point is a Layer 3 switch or router.

Interlocking Across Corporate Backbones

Interlocking across corporate backbone systems should be only considered by expert users. There are baseline requirements for the backbone that must be met, such as multicast routing protocols, spanning VLANs, redundancy protocol interaction, etc. that must be extensively tested, and timing characteristics must be confirmed before deployment. The impact on control system throughput (determinism, control data latency) must be extensively tested and evaluated for a successful deployment.

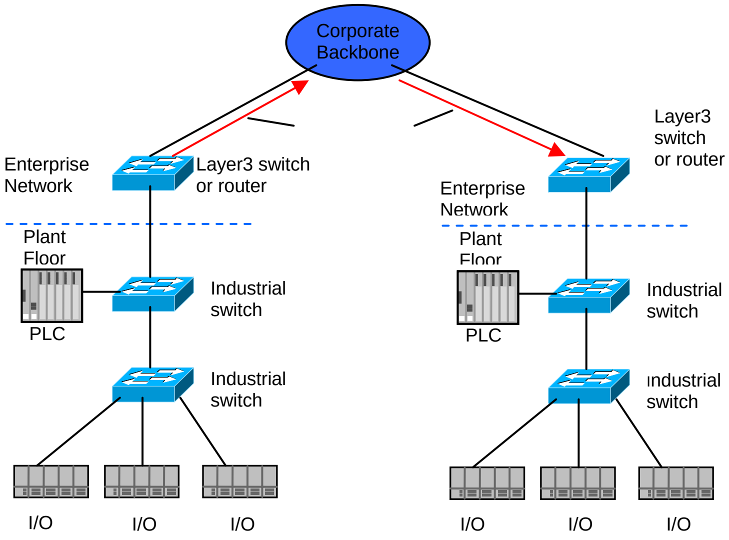

Figure 6-12 Example of an Interlocked Integrated Network.

6.5 Application Summary

Characterizing network traffic on EtherNet/IP helps determine which type of network will provide the best balance of performance and installed cost. The recommendations provided here are aimed at optimizing network—and ultimately control system—performance. They are based on the use of switches with the increasing levels of functionality needed to handle more devices with more traffic and short, implicit I/O control response times. Routers are used to selectively link these systems with enterprise or plant networks without disrupting any of the networks. Table 6-1 summarizes the relative degree that the application considerations discussed here are factors in the four general system types.

Table 6-1 Application Factors for System Types

	System Type			
Application Factors	**Isolated System Single Controller**	**Isolated System Multiple Controllers**	**Enterprise Network Connected**	**Enterprise Network Integrated**
Isolation of control system vs. plant network traffic and prevention of multicast traffic overload of control devices	Use of lower-cost unmanaged switches possible.	Requires managed (IGMP snooping) switches.	PC with 2 NICs possible; use of a CIP gateway preferred.	Requires managed switches with IGMP snooping and VLAN functionality.
	Added traffic design calculations if unmanaged switches are used.	VLAN functionality may optionally be needed.	Requires managed (IGMP snooping) switches.	Will likely require Layer 3 switches or routers at the join points.
			VLAN functionality may optionally be needed.	

Appendix A Recommendations for Ethernet Switches in EtherNet/IP Systems summarizes the required and desirable switch features for the various system types described here.

7 Selecting Components

This detailed look at the component features behind successful EtherNet/IP networks can be used as a checklist to help ensure the chosen components can provide the right functionality and withstand their industrial environments. Included here are connectors, wire and fiber-optic cabling, switches, routers, and equipment housings.

7.1 Sealed Connectors: Key to Successful Operation

Standard RJ-45 connectors, also called 8-way modular connectors, are not designed to operate in many industrial environments. Lower Common Mode Rejection Ratio (CMRR) cable types, also called balanced cables, may not provide adequate protection from industrial electrical noise levels and the potentially high amount of cross-talk. In high-vibration applications, additional mechanical securing means may also be required. For "outside the panel" connections, commercial-off-the-shelf RJ-45 connectors are not sealed to meet the IP65 and IP67 specifications for EtherNet/IP. Therefore, *The EtherNet/IP™ Specification*[22] includes an industrial version of the RJ-45 connector rated to meet IP65 and IP67. An M12 4-pin "D"-coded connector also is included to provide a smaller connector form factor.

Connectors play an important role in any EtherNet/IP network, and if they are going to be used in harsh industrial environments, they need to be sealed to keep out moisture and chemicals, and they have to be able to withstand high vibration. While EtherNet/IP uses standard RJ-45 connector technology, about the only locations in which these connectors can be used is inside sealed equipment cabinets where they are protected from plant-floor environments or in the relatively safe environments of control rooms.

7.1.1 Approved Sealed Version of RJ-45 Connectors

If an application requires cable and connectors to be exposed to liquid or dust contaminates, then the connectors need to be protected from chemicals, dust, solvents, oil, etc. RJ-45 (8-way modular) connectors and bulkheads sealed to meet IP65 and IP67 standards ratings are used to protect the wiring and plug connector pins from these contaminants. The IP65 and IP67 ratings specify that the connectors can withstand immersion in water at a depth of 1 meter (3.28 ft) for 30 minutes without incursion. IP65 and IP67 sealed RJ-45 (8 way modular) connectors are shown in Figure 7-1 and Figure 7-2 and have the following features:

- Bayonet-style flange
- Standard RJ-45 jack
- Seal at face of jack
- Accommodation for PCB mount, flange mount, or bulkhead mount

[22] ODVA, The CIP Networks Library, Vol. 1 and 2, *The EtherNet/IP™ Specification.* Ann Arbor: ODVA, Inc., 2006. www.odva.org. CD-ROM.

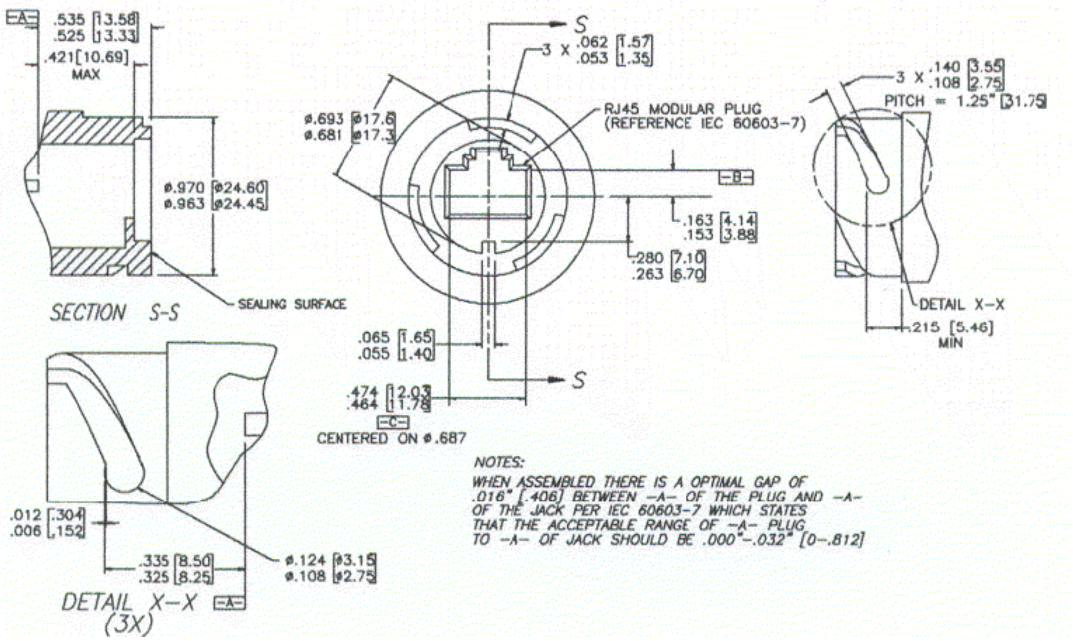

NOTES:
WHEN ASSEMBLED THERE IS A OPTIMAL GAP OF .016" [.406] BETWEEN -A- OF THE PLUG AND -A- OF THE JACK PER IEC 60603-7 WHICH STATES THAT THE ACCEPTABLE RANGE OF -A- PLUG TO -A- OF JACK SHOULD BE .000"-.032" [0-.812]

Figure 7-1 RJ-45 Sealed Connector.

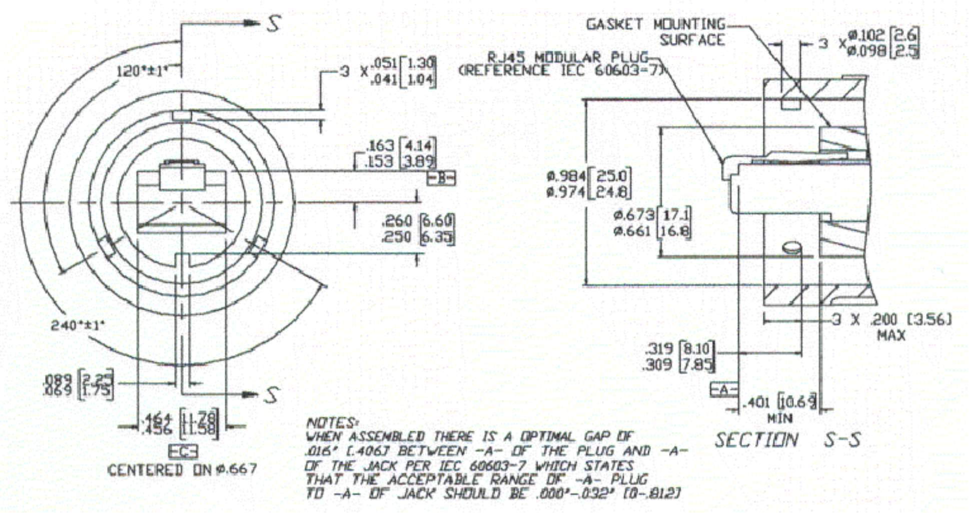

NOTES:
WHEN ASSEMBLED THERE IS A OPTIMAL GAP OF .016" [.406] BETWEEN -A- OF THE PLUG AND -A- OF THE JACK PER IEC 60603-7 WHICH STATES THAT THE ACCEPTABLE RANGE OF -A- PLUG TO -A- OF JACK SHOULD BE .000"-.032" [0-.812]

Figure 7-2 RJ-45 Sealed Plug.

7.1.2 M12 “D”-Coded Connector for EtherNet/IP

The EtherNet/IP™ Specification also allows the use of the four-pin, “D”-coded M12 connector for IP65- and IP67-rated environments. This device is shown on the left in Figure 7-3 and is the pin-and-socket style connector for EtherNet/IP. It will be familiar to many engineers from its standard version, which is in widespread use in industrial applications, such as with sensors and actuators. The connectors are not quite the same, however, and for good reason. The “D-coding” keying is different from the standard M12 connector to prevent accidental connection of other devices that may have high enough voltages to damage Ethernet circuits. The connector provides for two-pair shielded cables, and a variety of cables can be used to withstand most industrial conditions.

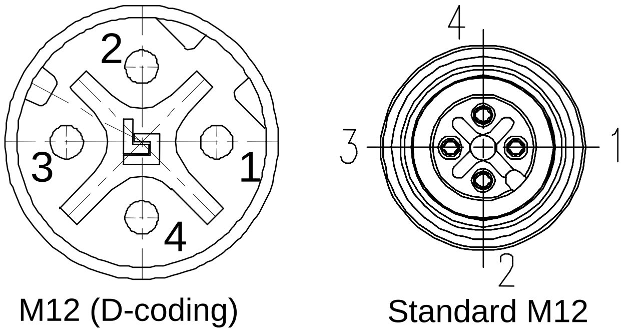

Figure 7-3 Comparison of “D”-Coded and Standard M12 Connectors.

7.1.3 Fiber-Optic Connectors

For fiber optics, there are three types of approved connectors: the quarter-turn LC, ST and the SC connectors. The ST (for straight-through) is shown on the left in Figure 7-4, and the SC (for standard connector) is shown on the right. The LC connector is a smaller version of the SC. A pair of fibers are used for each link to support separate receive and transmit paths. Therefore, like twisted-pair connections, receivers must be tied to transmitters. Ports on both the DTE and DCE are marked as TX (transmit) and RX (receive) to guide connections.

Figure 7-4 ST (left) and SC (right) Fiber-Optic Connectors.

7.2 Cabling Considerations: Twisted-Pair vs. Fiber-Optic Cable

Deciding whether to use fiber-optic cable or twisted-pair copper cable depends on specific application and environmental requirements. Users need to compare capabilities to find the right match for their requirements.

- Fiber-optic cable is inherently impervious to noise from electrical and magnetic fields.
- Fiber-optic cable may cost more per foot than standard, twisted-pair cable; however, users should consider the cost of installation for either type of cable.
- Fiber-optic cable has fairly stringent requirements regarding bend radiuses. (See *7.2.5 Bend Radius / Pull Force of Fiber*.)
- Shielded twisted-pair cable has been found to be resistant to some types of electrical noise; however, ground loops must be avoided. (See *7.2.1 When to Use Shielded Twisted-Pair Cabling*.)
- Fiber-optic cable can provide immunity to lightning and is a good choice between buildings or in other outdoor areas.
- Fiber-optic cable is a good option within hazardous areas. If it is necessary to mix communications cabling with high-voltage runs, fiber optics is superior to twisted-pair cabling due to the limited breakdown voltage rating of twisted-pair cable and potential noise coupling.
- Although not completely immune to security breaches, fiber-optic cable is much more difficult to tap into and access unauthorized data.
- Twisted-pair runs are limited to 100 meters (328 ft); fiber-optic cable can provide much longer transmission distances. Typical multi-mode installations allow distances up to 2 kilometers (1.24 mi), but single-mode installations can be 30 kilometers (18.64 mi) or greater.
- With both fiber-optic and copper cables, the cable jacket specifications for weld splatter, moisture, and UV (sunlight) resistance must be considered.

7.2.1 When to Use Shielded Twisted-Pair Cabling

Engineers need to select the correct cable for the environment. Cable selection is dependent on the planned cable route. Shielded cable may perform better than unshielded cable in high-noise industrial environments. In particular, if an application is in a high-noise environment or a cable must be run in close proximity to noise-radiating sources, shielded or fiber-optic cables should be used.

Shielded cables are appropriate if an application includes one or more of the following:

- Motor-control centers
- Induction-welding processes
- Proximity to high-power RF radiation

- Electrostatic processes
- High-current devices (greater than 100 amps)

7.2.2 Terminating Copper Cable

Termination of copper cables is a simple process. Copper cable termination tools and supplies are available from many convenient sources. Figure 7-5 shows a typical RJ-45 Ethernet connector. Studies have shown, however, when communicating at 100 Mbps, poor cable terminations can increase the number of retransmissions. This will reduce system throughput.

Attention also must be paid to the selection of wire pairs used when terminating a copper Ethernet cable. The proper wire pairs as defined in Table 7-1 must be used to preserve the cable's ability to reject noise and achieve the proper impedance.

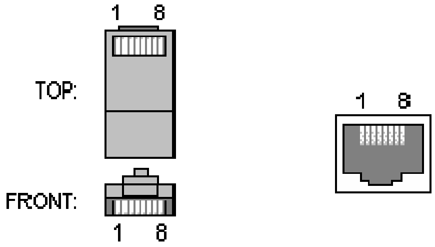

Figure 7-5 Views of an RJ-45 Ethernet Connector.

Table 7-1 Pinouts for RJ-45 Ethernet Connectors

Pin No.	Function	Color EIA/TIA 568B	Color EIA/TIA 568A
1	TX +	White/Orange	White/Green
2	TX -	Orange	Green
3	RX +	White/Green	White/Orange
4		Blue	Blue
5		White/Blue	White/Blue
6	RX -	Green	Orange
7		White/Brown	White/Brown
8		Brown	Brown

7.2.3 When to Use Fiber-Optic Cabling

Fiber should be considered when the required length is greater than 100 meters (328 ft) between two devices. If an application creates high magnetic and electric fields, fiber should be considered. Applications such as magnetic galvanizing annealing furnaces and high-power RF cross-linking machines are ideally suited to the use of fiber-optic cabling.

7.2.4 Types of Optical Fiber

Optical fiber consists of five basic elements: core, cladding, strength members, coating, and outer jacket. The core, constructed of either glass or plastic, provides the basic means for transmitting the light energy down the cable. The cladding prevents the light from exiting the core and being absorbed by the cable itself. The coating protects the core while providing strength. An overall jacket provides final protection and may consist of other strengthening and protective elements. Strength members provide protection when pulling the cable and for securing the connector.

Classification of Optical Fibers

Optical fibers are classified by their diameter in microns. A micron is a millionth of a meter and is designated by the symbol µm. The core, cladding, and coating frequently are specified using slashes (/) to separate the values. For example, 50/125/250 means that the core is 50 µm, the cladding is 125 µm, and the coating is 250 µm. Since the cable dimensions all pertain to the concentric diameters of the various elements, a concise way of specifying the fiber is to list only the core and cladding sizes. In the above example, this fiber would be classified as 50/125. Core sizes range from as small as 5 µm to as large as 1000 µm.

Depending on the core size, the cable transmits light in either one or two modes: single-mode or multi-mode. For each fiber-optic link, a pair of fibers is required. Usually several fibers are pulled to provide spares. The most popular multi-mode fiber-optic cables are 62.5/125 µm and 50/125 µm. A single-mode fiber-optic cable will have an inner core diameter of 10 µm or less.

7.2.5 Bend Radius / Pull Force of Fiber

Bend radius is defined as the radius of curvature that a fiber-optic cable can bend without causing any adverse effects. Each manufacturer will have different ratings for each of its products. It’s important to refer to the fiber-optic cable manufacturer’s data sheet for specific pulling force and bend radius recommendations when installing fiber-optic cable.

In the absence of this data, the pulling force should be limited to no more than 222 N (50 lbs) and the bend radius under load generally should be limited to 20 times the outside diameter of the cable and to 10 times the outside diameter under static conditions. For more information on fiber-optic bend radius, see the BICSI (a telecommunications association) *Telecommunications Distribution Methods Manual*.[23]

[23] BICSI, Inc., *Telecommunications Distribution Methods Manual*, 11th Edition, BICSI, Inc., Tampa: 2006. www.bicsi.org

7.2.6 Terminating Fiber-Optic Cable

Proper termination of fiber-optic cable can be a difficult process. If not performed properly, poor connections can reduce the distance capability of the system. Tools are available to help judge the quality of the termination and the fiber-optic cables.

Fiber connectors should only be installed by trained personnel. The channel performance can be severely degraded by improper connector installation.

In some cases, field splices may be required. Fusion splices should always be considered over a mechanical type of splice. Fusion splices should be performed by trained installers.

7.3 Key Switch Selection Issues

The additional functions offered by managed switches can facilitate faster start-ups, allow more devices to be connected without overloading the network, and add other capabilities often attributed to routers.

Managed infrastructure devices, such as switches and routers, are “intelligent” plant-floor devices. Users of managed infrastructure devices can enable or disable different functions—just like PLCs. Downloading the configuration parameters of a managed device, such as for a replacement switch, is no different from reloading a replacement PLC with the correct program. Also like PLCs, if multiple switches have different configurations, the engineer must maintain a library of configurations for specific switches. Table 7-2 summarizes the types of functions that may be found in a managed versus an unmanaged switch.

Table 7-2 Managed Switch versus Unmanaged Switch Functions

Switch Function	Managed Switch	Unmanaged Switch
Switched Ethernet: Eliminate collisions	Yes	Yes
10 Mbps to 100 Mbps data rates: Auto negotiation	Yes	Yes
10 Mbps to 100 Mbps data rates: Manual setting	Optional	Not available
Multicast message handling	Optional	Not available
Quality Of Service options	Optional	Optional
Security options	Optional	Not available
Virtual LAN: Forward messages	Yes	Yes
Virtual LAN: Configure/assign devices to a VLAN	Optional	Not available
Diagnostics: LED-based	Yes	Yes
Diagnostics: Software accessed (remote monitoring, web-based, set trap addresses)	Optional	Not available

7.3.1 Address Table Size Limits Control Device Quantity

Switches learn what devices are connected to each port by noting the address of each device that transmits (is the source of) a message. These addresses are stored in a look-up table in the switch's memory. Each switch has a finite capacity for stored addresses—usually several thousand. A switch updates its list using a process termed aging by periodically (ranging from tenths of seconds to minutes) removing addresses from which it has not received a message. If it receives a subsequent message from the device, the switch restores the device address to the table.

If a factory-floor system has thousands of connected devices, and it interconnects other systems that regularly broadcast messages that may pass through part of the control system, it is possible for a switch to run out of address table memory. Routers can block undesired broadcasts. Undesired broadcasts can also be contained through use of VLANs.

7.3.2 Auto-negotiation vs. Manual Setting of Switched Port Data Rates

The ports on most switches support both 10 Mbps and 100 Mbps—and many, 1 Gbps—data rates and half- or full-duplex operation. The capability of switches to automatically configure to the data rate (and half- versus full-duplex operation) of another auto-negotiation device that is connected can be a drawback in certain cases. When two auto-configuration devices are connected, they will first attempt to connect at the highest data rate, usually 100 Mbps. This can be undesirable in two cases:

- High electrical noise environments with twisted-pair cable. Manually setting the ports to 10 Mbps can provide a higher degree of noise immunity when using twisted-pair cable.

- Fiber-optic cable runs. Auto-negotiation is only offered for twisted-pair cabling and 10-Mbps fiber segments. The 100 Mbps data rates for fiber cables must be set manually.

7.3.3 Multicast Messaging Handling Option

EtherNet/IP can use multicast Ethernet frames for real-time communications. A typical switch will receive multicast (and broadcast) messages, and then rebroadcast them to the remaining switch's ports. Layer 2 and Layer 3 switches with IGMP snooping or equivalent functions learn which multicast messages to send to which ports.

7.3.4 Quality of Service Options

Switches can prioritize traffic based on MAC addresses, IP addresses, switch port numbers, and protocol identification information. This ability to prioritize and provide a different QoS based on the message type can help improve real-time performance. It requires the user to enter QoS configuration information into the Layer 2 or Layer 3 switch.

7.3.5 Security Options

Switches can maintain port-based Access Control Lists (ACL). This capability prevents unauthorized access into the control network based on the user entering a MAC address or a port-number-based ACL. This approach is also used to disable unused ports, so they cannot be used for unauthorized network access.

7.3.6 Broadcast or Multicast Threshold Filters

Switches may have a fixed or user-variable option for restricting the maximum allowed level of broadcast or multicast traffic. Once the threshold is exceeded, it may turn off the port until the traffic falls below the threshold level or shut down the port completely based on the vendor's implementation. This feature can help lessen the effects of abnormal traffic conditions. For high-speed control applications, the threshold level must be checked to confirm that it does not interfere with normal traffic.

7.3.7 Additional Functions

The switch functions described above provide an overview of basic capabilities. *Appendix A Recommendations for Ethernet Switches in EtherNet/IP Systems* contains descriptions of additional industrial switch functions, and how they relate to the various types of application scenarios described in *6 Infrastructure Application Scenarios*.

7.4 Distributed Enclosures and Environmental Considerations

Commercial switches, routers, and other equipment are fine for the office closet. If, however, commercial network equipment is to be located on the plant floor, it must be housed in industrial enclosures that are capable of protecting it from plant-floor environments. Industrial-grade equipment is the better option as it is designed to withstand harsh conditions. In either case, environmental factors can be critical to reliable operation, and all equipment should be protected accordingly. Lost or corrupted data is bad enough, but harsh conditions—especially heat and vibration—can slowly destroy unprotected equipment.

7.4.1 Environmental Specification Considerations

The required specifications for Ethernet devices used in industrial applications are dependent on where and how the devices are to be mounted. DIN-rail versus rack mounting and the use of 120 Vac versus 24 Vdc power need to be considered. Some additional considerations follow.

7.4.2 Temperature

Industrial manufacturing facilities can contain a wide range of environments, even within the same facility. The environment outside the enclosure, as well as heat-generating devices within the enclosure, needs to be considered. Without added environmental conditioning equipment, traditional (commercial) infrastructure components with 40^{o} C to 45^{o} C (104^{o} F to 113^{o} F) ratings should be used in control rooms or lighter industrial environments away from heat-producing machines or processes.

Cabinet fans or air conditioners may be needed for hot environments or where heat-producing PCs or motion control devices are mounted in the same enclosure. When mounted near or on heat-producing machines or processes, and where space, maintenance or other considerations prevent the use of air conditioners or fans, infrastructure components with PLC ratings of 55^{o} C (131^{o} F) or higher are available. Operating humidity specifications of 85% to 95% RH (non-condensing) should be considered for processing applications where water is present (wood, paper, food, etc.) and in non-air-conditioned plants in high-humidity locations. In corrosive atmospheres, electronics should be sealed completely to protect circuits from caustic or acidic vapors that can strip away conductors on printed circuits.

7.4.3 Electrical noise

Electrical noise is produced by motors, welders, and power-switching control devices, such as drives and servo controls, RF wireless transmitters, incoming power lines, and a variety of other sources. The higher data rates of Ethernet, especially fast Ethernet (100 Mbps), require greater attention to cabling issues and equipment noise specifications. Equipment rated under the following standards should be considered for electrically noisy environments:

- IEC 61000 (electromagnetic interference)
- IEC 61000-4-2 (electrostatic discharge)
- IEC 61000-4-3 (radiated noise immunity)
- IEC 61000-4-4 (noise burst)
- IEC 61000-4-5 (surge)
- IEC 61000-4-6 (conducted noise immunity)
- IEC 61000-4-8 (magnetic field immunity)
- IEC 61000-4-11 (voltage dips and interruptions)

It is important to remember when comparing specifications, that many of the above mentioned standards specify different immunity levels or criteria ratings. Vendors are a good source for any additional criteria that may apply.

7.4.4 Shock and Vibration

When mounting infrastructure components on or near vibrating machinery, such as punch or stamping presses, etc., it is important to carefully review all shock and vibration specifications for all the components. Equipment rated under the following standards should be considered for higher shock and vibration environments:

- IEC 61131-2, IEC 60068-2-27 (shock)
- IEC 60068-2-6 (vibration)

7.4.4.1 Certifications

Typical certifications for industrial applications include CE, EN 60950, UL/cUL1950, and UL/cUL 508. For use in potentially explosive environments in the United States, added certifications such as UL/cUL 1604 Class 1, Div. 2 (A, B, C, D) or FM 3611 Class 1, Div. 2 may be required. In Canada or Europe, certification to a zone rating system may be required.

Appendix A Recommendations for Ethernet Switches in EtherNet/IP Systems

In industrial applications, infrastructure devices are mounted in the same enclosures as the control devices, as opposed to the dedicated switch cabinets used in the office world. Therefore, the infrastructure devices used in an EtherNet/IP system must withstand the same environmental conditions as the control devices and be compatible with their voltage and mounting conventions.

Personnel for System Start-Up and Maintenance

While a facility may have an Information Technology Department staffed to support hundreds of PCs, the same facility can have over a thousand control devices with the potential for Ethernet connectivity. The basic startup, maintenance diagnostics, and configuration of industrial infrastructure devices must be supported by existing plant and control engineering staffs. Equally important, the infrastructure devices need to be compatible with IT tools and network standards.

The recommendations presented here are consistent with these objectives:

- Consider the impact of system cost while providing switch functionality adequate for successful initial I/O control applications and expansions.

- Use only industry standards (i.e., IEEE, IEC, TIA, etc.)—not company-owned proprietary functions.

- Include common functions available from three or more industrial infrastructure vendors.

- Use the application-based system types described in *6 Infrastructure Application Scenarios* to accommodate EtherNet/IP applications that range from small initial applications to large enterprise-integrated systems.

This appendix categorizes certain switch features as REQUIRED, RECOMMENDED, or APPLICATION-SPECIFIC.

- REQUIRED = functions that are the minimum required for typical performance in the majority of installations of the application type described in *6 Infrastructure Application Scenarios*. When in doubt, use the next higher application section’s recommendations. In addition to performance factors, the most likely required plant-floor installation and maintenance functions are included. Each section’s requirements build on the previous section’s requirements.

- RECOMMENDED = functions that are not necessarily needed for all application scenarios of this type, but typically improve or simplify maintenance or provide added performance in certain situations.

- APPLICATION SPECIFIC = functions that may be significant based on the specific application, but not necessarily for the particular type of system architecture (scenario). These functions typically cross several application categories. The more features a switch has, the broader the application range, and potentially, the higher the cost.

Isolated Systems: Small Scale

Application Range: Isolated control network with a single controller and a low device count

REQUIRED:

1. Speeds are 10 Mbps through 100 Mbps with full-duplex operation on all ports. Full-duplex operation eliminates collisions, a key requirement for deterministic control.

2. LED diagnostics have a Link LED per port (a physical cable connection is acceptable) and a network traffic activity indication per port. The traffic indication allows the user to differentiate between no traffic, intermittent traffic, and high traffic.

RECOMMENDED:

3. The switch auto-negotiates the data rate, the duplex settings, and provides auto-crossover (MDI/MDX) functionality to reduce device set-up requirements and cable compatibility problems.

4. The switch supports a maximum operating temperature of 55° C (131° F) or higher.

5. The switch meets the IEC Std 61000-4 or equivalent standards for noise immunity and emission.

Application Notes:
The primary use of these types of switches is for small applications (less than10 devices) with a small number of CIP connections and slower RPIs, or applications that are lower cost extensions to IGMP-enabled switches.

The above specifications relate to switches without IGMP snooping (unmanaged switches). Due to the multicast operation of EtherNet/IP, one must either calculate the total multicast traffic to which each connected unmanaged switch will be exposed or test the configuration in advance. Failure to do this may result in devices overloaded with unwanted multicast traffic, which may cause communications to stop.

Isolated Systems: General Use

Application Range: Isolated control network with a single controller and a high (more than 10) device count, or an isolated network with multiple, interlocked controllers

REQUIRED:

1. All specifications included in *Isolated Systems: Small Scale*.

2. IGMP snooping: This provides for proper management of IP multicast traffic. More specifically, it provides for fast consumer joins and leaves of multicast streams (reducing memory use) and limits bandwidth-intensive IP multicast traffic to only those ports that are intended to consume it.

Lack of IGMP snooping function may result in flooding of the subnet with multicast traffic, potentially overloading devices on the subnet. This is because the default operation of switches treats multicast packets the same as broadcast packets, while much I/O messaging with EtherNet/IP uses IP multicast to distribute I/O data, which is consistent with the CIP producer-consumer model.

3. IGMP query: Each application must have at least one switch or router that can generate the periodic IGMP query commands used to determine which devices are members of which multicast groups.

RECOMMENDED:

4. Manual port settings: The switch should allow manual configuration of both port mode (full- or half-duplex) and speed selection (10 Mbps, 100 Mbps, or 1000 Mbps). Legacy devices may have auto-negotiation compatibility problems. In addition, in some exceptionally high electrical noise environments, manually reducing the 100-Mbps data rate to 10 Mbps provides higher noise immunity.

5. Web- and SNMP-accessible port status and diagnostics: A managed switch, which has its own IP address, should support SNMP and web server as means for remote monitoring and configuration. As switches are applied in larger scale or faster, time-critical applications, the need for diagnostics increases. Web interfaces allow diagnostic access by plant-floor personnel. SNMP diagnostics are needed to interface with network management (plant-floor or IT Department-based) and IT Department support tools. Port status information, typically a web page, allows fast determination of which ports are connected and their operating status. This is often required for basic cabling troubleshooting. RMON diagnostics provide statistics regarding the types and sizes of packets being received by any port, as well as port utilization. This provides basic network traffic statistics to optimize response times and locate noise-induced disruptions without a sophisticated traffic analyzer.

6. Port mirroring: This function allows traffic from any port on a switch to be duplicated on a spare port for advanced network diagnostic monitoring and troubleshooting. Port mirroring is the function that provides compatibility with IT-level personnel and network troubleshooting tools.

Application Notes:

The primary use of these types of switches is, in general, isolated system applications where all the devices need to interoperate within a single system. The use of IGMP snooping reduces traffic and allows larger systems and faster response times. While low-cost unmanaged switches may be used as an expansion to unmanaged switches, precautions must be taken to prevent the multicast overload of devices. For the highest performance systems with the fewest start-up or future expansion concerns, it is recommended that all switches have IGMP snooping functionality.

When connecting multiple controllers over EtherNet/IP, if controller-to-controller communications is needed (i.e., for interlocking or other peer-to-peer communications) on the same network, then IGMP snooping is required. If the multiple controllers are independent systems (no peer-to-peer communications), then they are either a collection of general, small isolated systems, or they connect to a supervisory network and are classified as an enterprise-connected system.

Large-scale Control and Enterprise Networking

Application Range: Enterprise-connected control system or enterprise-integrated control system

REQUIRED:

1. All specifications included in *Isolated Systems: General Use*.

2. Web- or SNMP-accessible port utilization and bandwidth diagnostics: In systems where enterprise connectivity and real-time control performance needs must be balanced, the ability to see the percent of bandwidth utilization of each port allows a faster means of locating traffic overload points. The RMON diagnostics can then be used to further dissect the nature of the high traffic, and corrective action can be taken by changing VLAN, routing, or other parameters.

RECOMMENDED:

3. For enterprise-connected applications without gateways, the use of industrial switches with VLAN capability is recommended. In this case, existing office (non-industrial) routers may be used to manage the selected data transfer between the control system and the plant network systems.

4. For enterprise-integrated applications, additional Layer 3 switches or routers are recommended to provide the real-time transfer of data while isolating disruptive traffic from the control system and enterprise networks.

Application Notes:
See *6.4 Enterprise-Connected and Integrated Control Systems.*

APPLICATION-SPECIFIC OPTIONS:
The functions noted here provide added performance or ease of maintenance. The actual need for these and any associated cost versus functionality tradeoffs need to be evaluated for each application. Some of the functions are still emerging in the industrial application range and may not yet have widespread multi-vendor support. Generally speaking, the more of these functions a given switch has, the greater the range of EtherNet/IP applications that can be addressed. Unless otherwise noted, the functions noted here relate to applications using managed switches (isolated, general use systems and higher).

System Performance:

1. IEEE Std 1588 functionality for motion control applications coordinated with CIP Sync: In order to implement CIP Sync applications, switches must support IEEE 1588 functionality. See *6.4.1 Networks Connected via CIP Gateways* and *The EtherNet/IP Specification*[24] for more information.

2. IEEE-based redundancy: IEEE Std 802.1D (Spanning Tree) and IEEE Std 802.1w (Rapid Spanning Tree) provide redundant backbone connections and loop-free networks. These previously separate standards have now been combined under the single IEEE Std 802.1D.

[24] ODVA, The CIP Networks Library, Vol. 1 and 2, *The EtherNet/IP™ Specification.* Ann Arbor: ODVA, Inc., 2006. www.odva.org. CD-ROM.

This simplifies network configuration and improves fault tolerance. Engineers can enable or disable the protocol as needed. Both of these standards support ring and tree/mesh topologies, integrate with existing office switches, and are well known by IT personnel. Network recovery times range from a few seconds for the Rapid Spanning Tree Protocol, to 30 seconds to 60 seconds for Spanning Tree Protocol.

3. IEEE Std 802.1p frame prioritization (QoS): This function allows control traffic to have a higher priority (serviced first by the switch) then non-time critical traffic. There are two basic levels of implementation common in industrial switches:

    - *Industrial switch recognition of pre-tagged packets:* Common in managed switches, emerging in unmanaged switches, this level of functionality is where the switch receives, reads, and places the packets in two or more (eight maximum) priority queues. The packets are prioritized via a managed switch or a router located in the system.

    - *Industrial switch both assigns tags and recognizes pre-tagged packets:* The managed switch typically assigns priorities by the port to which a particular device is connected. Routers can also assign priorities by device MAC or IP address. Other, more sophisticated options may include the ability to also set a CoS (Class of Service) that is configurable for I/O traffic (i.e., UDP, port 0x8AE) on EtherNet/IP and the support of configurable algorithms to manage multiple priority levels.

4. Port aggregation: In multi-switch applications, port aggregation allows traffic between switches to be shared between two or more physical connections, essentially increasing the bandwidth of the backbone connections. In this case, two 100-Mbps ports can be combined using IEEE Std 802.1ad link aggregation protocol to form a single 200-Mbps connection that supports load balancing.

5. One Gigabit or higher port data rates: For large-scale, enterprise-connected and enterprise-integrated systems, larger 1-Gbps backbone capacity may be desired. These connections are used for switch-to-switch uplink or backbone connections, not control connections to EtherNet/IP devices. EtherNet/IP industrial control devices support 10 Mbps and 100 Mbps data rates. The decentralized architecture of industrial Ethernet reduces the need for large concentrations of devices and higher than 100-Mbps connections.

6. Broadcast or multicast threshold filters: Switches may have a fixed or user-variable option for restricting the maximum allowed level of broadcast or multicast traffic. Once the threshold is exceeded, the port may turn itself off until the traffic falls below the threshold level or shut down the port completely based on the vendor’s implementation. This feature can help lessen the effects of abnormal traffic conditions. For high-speed control applications, the threshold level must be checked to confirm it does not interfere with normal traffic. Vendor-specific operation of this function, such as operation on ingress/egress, types of messages filtered (unicast, multicast, or broadcast), and action of the ports when the limit is reached (clipping versus a shutdown that requires manual reset) must be considered. Attention to the overall network design will reduce the need for these filtering functions.

Industrial System Maintainability

In addition to infrastructure-embedded functions, the use of plant-floor network management software provides added plant maintainability, and compliments the IT-based network management software already installed in many facilities. This type of software provides a simplified “PLC look” interface to the entire EtherNet/IP industrial system. The system-wide view allows quick identification of non-communicating devices, the centralized viewing of SNMP diagnostic “trap” messages, and the ability to “zoom” into detailed device-specific diagnostics using network management screens or the device’s embedded web pages. These packages also allow easy upload/download (via open Trivial File Transfer Protocol [TFTP]) of switch/device configurations and upgrades. Because these packages use the industry standards, such as SNMP protocol, network management software from one vendor can work with other vendor’s switches.

The functions below relate to infrastructure (switch or router) embedded functions.

1. Alarm contact: The addition of a physical relay-out allows annunciation of power outages and/or disabling of key links (cable cut, device powered down, etc). The alarm contact is connected to either physical alarms, lights, etc., or to automation inputs (i.e., spare PLC inputs) for operator intervention. This provides added cabling and power supply diagnostics for non-IT or engineering trained personnel.

2. Serial configuration port for managed switches: A serial configuration port allows the IP addressing of managed switches or the confirmation of an IP address by plant personnel (using Hyper Terminal with a laptop) without special software, training, or network access.

3. Automatic re-addressing of replacement devices with DHCP Option 82: DHCP Option 82 allows dynamic IP address assignment for the attached nodes to either be fixed, based on the port they are plugged into (instead of their MAC_ID), or centrally managed (typically based in the IT Department) in a server that contains the master IP address list. This function eliminates the need to manually assign the same IP address to a replacement device’s new MAC address.

4. Complete web access to all switch parameters: If the EtherNet/IP system is to be primarily maintained by plant and control engineering personnel, there is a greater need to have all switch configuration parameters accessible from standard laptop browsers. Parameters that can only be accessed via SNMP network management software require the purchase of network management software or reliance on IT Department tools and personnel. If the control network is to be primarily maintained by IT personnel, this is not as great a requirement.

5. Port mirroring: This function allows traffic from any port on a switch to be duplicated on a spare port for advanced network diagnostic monitoring and troubleshooting. Port mirroring is the function that provides compatibility with IT-level network troubleshooting tools.

6. Physical storage/retrieval of switch configuration: As an alternative to network-based auto-addressing and auto-reconfiguration, memory cards or modules may be used to store a switch’s configuration. If replacement of a switch becomes necessary, plant personnel can remove the card or module from the previous switch, insert it into the new switch, and have

the configuration read-in automatically. This allows switch maintenance and replacement with non-engineering personnel or where access to switch configuration data is restricted.

7. Security functionality: Security appropriate for Ethernet-based control systems is currently being evaluated. Most managed switches support at least port disable functions. Current security measures that are becoming more common in managed switches include the following:

   - Port disable: The ability to disable/enable spare maintenance ports to prevent casual access.

   - Port security: The ability to enter a list of one or more authorized devices' MAC or IP addresses. Connected devices that do not match the list are prevented from communicating and typically also generate an SNMP alarm message.

   - VLANs: VLANs can be used to restrict traffic and network access. Unauthorized access in one area can be restricted to that section of the application instead of the entire EtherNet/IP application.

   - Radius authentication: This includes IEEE Std 802.1X for plant-floor based programming stations and PCs. It prevents the non-authorized laptop or PC that is new to the area from being attached to and used on the network without being authenticated by a central server.

   - SSL/SSH: These security measures allow the secure remote monitoring and configuration of devices via the web or command line (CLI) Telnet.

   - Firewalls and VPNs: These allow the secure movement of data across unsecured lines and keep access to sensitive areas secure from unauthorized personnel.

More information related to various encryption techniques will be added in future revisions of this document.

Summary

The minimum functional switch requirements related to EtherNet/IP vary based on the size, performance, functionality and network connectivity requirements of the application. (See *6 Infrastructure Application Scenarios* for more information.) For EtherNet/IP applications, these are summarized with three levels of required functionality:

- Isolated systems, small scale
- Isolated systems, general use
- Large-scale control and enterprise networking

In any specific application, there are additional optional functions that can increase performance and maintainability of the system. The following represent current widely available or quickly emerging functions.

APPLICATION-SPECIFIC FUNCTIONS

Overall System Performance

- IEEE Std 1588
- IEEE redundancy
- IEEE Std 802.1p frame prioritization
- Port trunking
- Routing
- One Gbps (and higher) data rates for the uplinks
- Broadcast or multicast threshold filters

Added Industrial System Maintainability

- Alarm contact
- Serial configuration port for managed switches
- Automatic re-addressing of replacement devices with DHCP Option 82
- Complete web access to all switch parameters
- Port mirroring
- Physical storage/retrieval of switch configuration
- Security functionality

Appendix B Overview of the OSI Model, EtherNet/IP and CIP

The Open System Interconnection (OSI) Reference Model describes how information from a software application in one computer moves through a network medium to a software application in another computer. The OSI Reference Model is a conceptual model composed of seven layers, each specifying particular network functions as shown in Figure B-1. The International Organization for Standardization (ISO) developed the model in 1984, which is now considered the primary architectural model for inter-computer communications.

The OSI Model divides the tasks involved in moving information between networked computers into seven smaller, more manageable task groups. A task or group of tasks is then assigned to each of the seven OSI layers. Each layer is reasonably self-contained so that the tasks assigned to each layer can be implemented independently. This enables the solutions offered by one layer to be updated without adversely affecting the other layers.

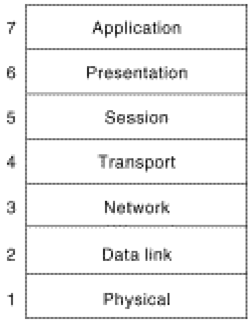

Figure B-1 The OSI Reference Model's Seven Independent Layers.

Characteristics of the OSI Layers

The seven layers of the OSI Reference Model can be divided into two categories: upper layers and lower layers as shown in Figure B-2.

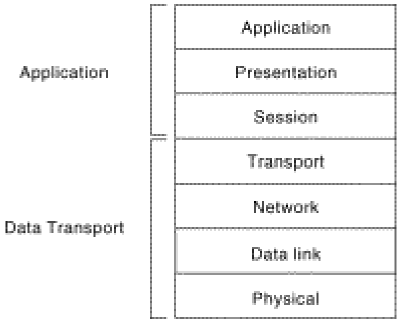

Figure B-2 The OSI Model's Two Sets of Layers.

The upper layers deal with application issues and generally are implemented only in software. The highest layer, the application layer, is closest to the end user. Both users and application layer processes interact with software applications that contain a communications component. The term “upper layer” is sometimes used to refer to any layer above another layer in the OSI Model.

The lower layers handle data transport issues. The physical layer and the data link layer are implemented in hardware and software. The lowest layer, the physical layer, is closest to the physical network medium (the network cabling, for example) and is responsible for actually placing information on the medium.

Communication between Systems

Information being transferred from a software application in one computer system to a software application in another must pass through the OSI layers. For example, if a software application in System A has information to transmit to a software application in System B, the application program in System A will pass its information to the application layer (Layer 7) of System A. The application layer then passes the information to the presentation layer (Layer 6), which relays the data to the session layer (Layer 5), and so on—down to the physical layer (Layer 1). At the physical layer, the information is placed on the physical network medium and is sent across the medium to System B.

The physical layer of System B removes the information from the physical medium, and then its physical layer passes the information up to the data link layer (Layer 2), which passes it to the network layer (Layer 3), and so on, until it reaches the application layer (Layer 7) of System B. Finally, the application layer of System B passes the information to the recipient application program to complete the communication process.

Interaction between OSI Model Layers

A given layer in the OSI model generally communicates with three other OSI layers: the layer directly above it, the layer directly below it, and its peer layer in other networked computer systems.

The data link layer in System A, for example, communicates with the network layer of System A, the physical layer of System A, and the data link layer in System B as shown in Figure B-3.

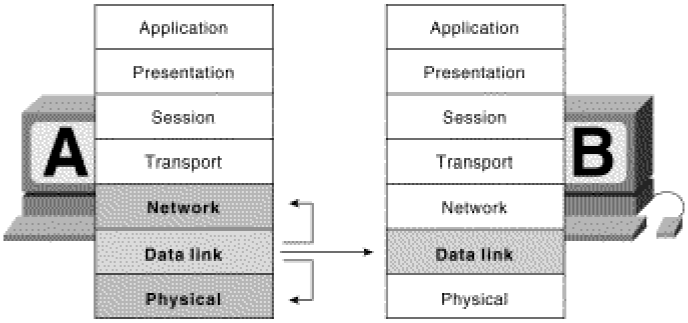

Figure B-3 How OSI Layers Communicate.

OSI Layer Services

One OSI layer communicates with another layer to make use of the services provided by the second layer. The services provided by adjacent layers help a given OSI layer communicate with its peer layer in other computer systems.

Basic Elements of Layer Services

Three basic elements are involved in layer services: the service user, the service provider, and the service access point (SAP).

In this context, the service user is the OSI layer that requests services from an adjacent OSI layer. The service provider is the OSI layer that provides services to service users. OSI layers can provide services to multiple service users. The service access point (SAP) is a conceptual location at which one OSI layer can request the services of another OSI layer. Figure B-4 shows how these three elements interact at the network and data link layers.

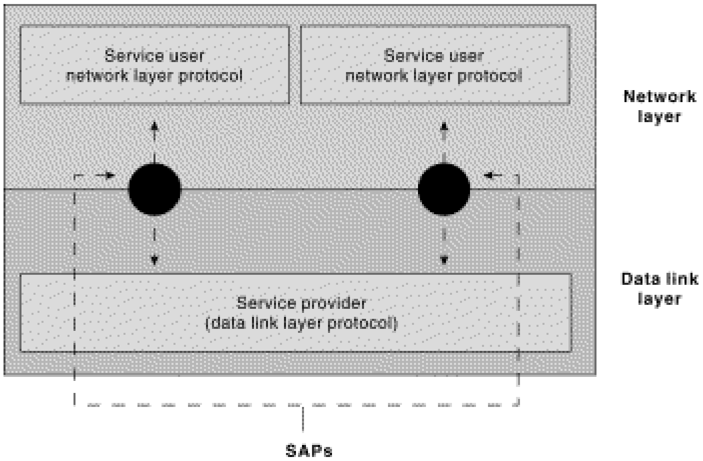

Figure B-4 Service Users, Providers, and SAPs Interact at the Network and Data Link Layers.

OSI Model Layers and Information Exchange

In one computer, the seven OSI layers use various forms of control information to communicate with their peer layers in other computer systems. This control information consists of specific requests and instructions that are exchanged between peer OSI layers.

Control information typically takes one of two forms: headers and trailers. Headers are prepended to data that has been passed down from upper layers. Trailers are appended to data that has been passed down from upper layers. An OSI layer is not required to attach a header or a trailer to data from upper layers.

Headers, trailers, and data are relative concepts, depending on the layer that analyzes the information unit. At the network layer, for example, an information unit consists of a Layer 3 header and data. At the data link layer, however, all the information passed down by the network layer (the Layer 3 header and the data) is treated as data.

In other words, the data portion of an information unit at a given OSI layer potentially can contain headers, trailers, and data from all the higher layers. This is known as encapsulation. Figure B-5 shows how the header and data from one layer are encapsulated into the header of the next lowest layer.

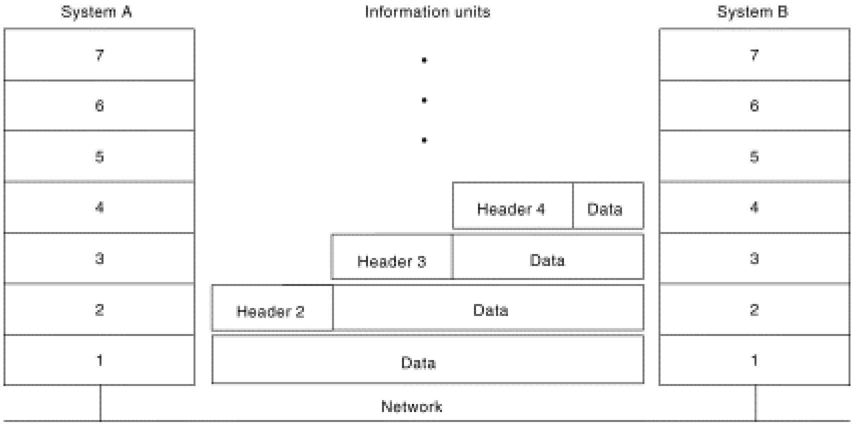

Figure B-5 Headers and Data Can Be Encapsulated during Information Exchange.

Information Exchange Process

The information exchange process occurs between peer OSI layers. Each layer in the source system adds control information to the data, and each layer in the destination system analyzes and removes the control information from that data.

If System A has data from a software application to send to System B, the data is passed to the application layer. The application layer in System A then communicates any control information required by the application layer in System B by prepending a header to the data. The resulting information unit (a header and the data) is passed to the presentation layer, which prepends its own header containing control information intended for the presentation layer in System B. The information unit grows in size as each layer prepends its own header (and, in some cases, a trailer) that contains control information to be used by its peer layer in System B. At the physical layer, the entire information unit is placed onto the network medium.

The physical layer in System B receives the information unit and passes it to the data link layer. The data link layer in System B then reads the control information contained in the header prepended by the data link layer in System A. The header is then removed, and the remainder of the information unit is passed to the network layer.

Each Layer Performs the Same Actions

The layer reads the header from its peer layer, strips it off, and passes the remaining information unit to the next highest layer. After the application layer performs these actions, the data is passed to the recipient software application in System B, in exactly the form in which it was transmitted by the application in System A.

Physical Layer

The physical layer defines the electrical, mechanical, procedural, and functional specifications for activating, maintaining, and deactivating the physical link between communicating network systems. Physical layer specifications define characteristics, such as voltage levels, timing of voltage changes, physical data rates, maximum transmission distances, and physical connectors. Physical layer implementations can be categorized as either LAN or WAN specifications. Figure B-6 illustrates some common LAN and WAN physical layer implementations.

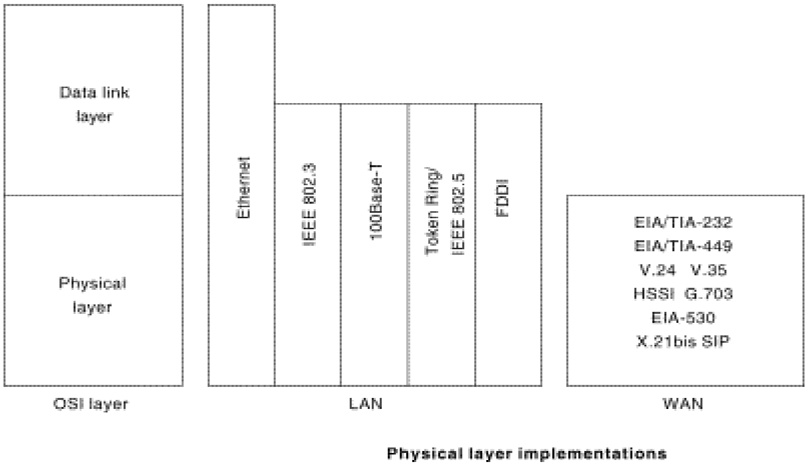

Figure B-6 Physical Layer Implementations Can Be LAN or WAN Specifications.

Data Link Layer

The data link layer provides reliable transit of data across a physical network link. Different data link layer specifications define different network and protocol characteristics, including physical addressing, network topology, error notification, sequencing of frames, and flow control. Physical addressing (as opposed to network addressing) defines how devices are addressed at the data link layer.

Network topology consists of the data link layer specifications that often define how devices are to be physically connected, such as in a bus or a ring topology. Error notification alerts upper-layer protocols that a transmission error has occurred, and the sequencing of data frames reorders frames that were transmitted out of sequence. Finally, flow control moderates the transmission of data such that the receiving device is not overwhelmed with more traffic than it can handle at one time.

IEEE has subdivided the data link layer into two sub-layers: Logical Link Control (LLC) and Media Access Control (MAC). Figure B-7 illustrates the IEEE sub-layers of the data link layer.

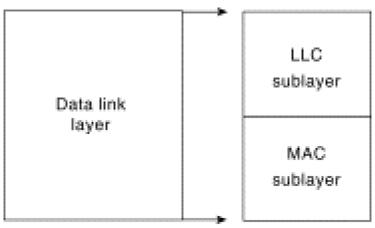

Figure B-7 The Data link layer Contains Two Sub-layers.

The Logical Link Control (LLC) sublayer of the data link layer manages communications between devices over a single link of a network. Defined in IEEE Std 802.2, LLC supports both connectionless and connection-oriented services used by higher-layer protocols. IEEE Std 802.2 defines a number of fields in data link layer frames that enable multiple higher-layer protocols to share a single physical data link.

The Media Access Control (MAC) sub-layer of the data link layer manages protocol access to the physical network medium. The IEEE MAC specification defines MAC addresses, which enable multiple devices to uniquely identify one another at the data link layer.

Network Layer

The network layer defines the network address, which differs from the MAC address. Some network layer implementations, such as the Internet Protocol (IP), define network addresses in a way that route selection can be determined systematically by comparing the source network address with the destination network address and applying the subnet mask. Because this layer defines the logical network layout, routers can use this layer to determine how to forward packets. Because of this, much of the design and configuration work for inter-networks happens at Layer 3, the network layer.

Transport Layer

The transport layer accepts data from the session layer and segments the data for transport across the network. Generally, the transport layer is responsible for making sure that the data is delivered error-free and in the proper sequence. Flow control generally occurs at the transport layer.

Flow control manages data transmission between devices so that the transmitting device does not send more data than the receiving device can process. Multiplexing enables data from several applications to be transmitted onto a single physical link. Virtual circuits are established, maintained, and terminated by the transport layer. Error checking involves creating various mechanisms for detecting transmission errors, while error recovery involves acting, such as requesting that data be retransmitted, to resolve any errors that occur. The transport protocols used on the Internet are TCP and UDP.

Appendix C EtherNet/IP and International Standards

EtherNet/IP is not just an open network based on the CIP specification. Its lower layers are based on Internet and Ethernet standards. In addition, *The EtherNet/IP™ Specification*[25] is rapidly being incorporated into international standards. This high degree of standardization assures users that EtherNet/IP systems will work with Commercial Off-The-Shelf (COTS) Ethernet infrastructure devices and Internet protocols, and be accepted in manufacturing systems around the world. Listed here are other international standards related to EtherNet/IP.

Internet and Ethernet Standards

The Internet Engineering Task Force (IETF) publishes its standards under a series called Request For Change documents (RFCs). EtherNet/IP is based on the following Ethernet and Internet standards:

RFC 768	User Datagram Protocol
RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 793	Transmission Control Protocol
RFC 826	Ethernet Address Resolution Protocol
RFC 894	Transmission of IP Datagrams over Ethernet Networks
RFC 1112	Host Extensions for IP Multicasting
RFC 2236	Internet Group Management Protocol, Version 2
ISO/IEC 8802-3:1996	Information Technology, Telecommunications and information exchange between systems, Local and metropolitan area networks, Specific requirements, Part 3: Carrier sense multiple access with collision detection (CSMA/CDD) access method and physical layer specifications.

Protocols and Profiles

EtherNet/IP is a part of the following IEC and CENELEC standards:

IEC 61158:2003 (Type 2)	Covers a wide range of application requirements from discrete manufacturing automation to process control. It includes the CIP application layer used by EtherNet/IP and other CIP-based networks, and the mapping protocol for putting CIP on TCP/IP/UDP
IEC 61784-1:2003 CP 2/2 Type 2	Includes the communication profile defining how to extract the relevant EtherNet/IP specifications from IEC 61158 documents
EN 61158:2003 (Type 2)	European equivalent of the corresponding IEC standard

[25] ODVA, The CIP Networks Library, Vol. 1 and 2, *The EtherNet/IP™ Specification.* Ann Arbor: ODVA, Inc., 2006. www.odva.org. CD-ROM.

EN 61784-1:2003 CP 2/2 Type 2	European equivalent of the corresponding IEC standard

Electronic Data Sheets (EDS)

The definitions contained in the Electronic Data Sheet (EDS) as defined in *The EtherNet/IP™ Specification* are in compliance with ISO 15745-Parts 1 and 4:2003. ISO 15745 is applicable to a wide range of application requirements from discrete manufacturing automation to process control. It defines an application integration framework, i.e., a set of elements and rules for describing integration models and application interoperability profiles.

Appendix D Glossary of Terms

Some of the terms defined below have been copied from [1] and [2] with some tailoring to EtherNet/IP terminology where appropriate.

Term	Definition
Address Resolution Protocol (ARP)	A protocol used to dynamically bind a high-level IP address to a low–level physical hardware address. Within the scope of this document it means a protocol that converts an IP address into an Ethernet address. ARP is used across a single physical network and is limited to networks, like Ethernet, which support broadcast.
BOOTstrap Protocol (BOOTP)	A protocol a node uses to obtain its IP Address, Subnet Mask and Gateway Address from a server.
broadcast	A transmission method, by which a packet is sent to multiple, unspecified recipients. Broadcast transmission is supported by Ethernet and IP protocols.
Dynamic Host Configuration Protocol (DHCP)	A protocol a node uses to obtain its IP Address, Subnet Mask and Gateway Address from a server. A superset of the BOOTP.
explicit messaging	Explicit messages can be sent as connected or unconnected messages. CIP defines an explicit messaging protocol that states the meaning of the message. This messaging protocol is contained in the message data. Explicit messages provide a one-time transport of a data item. Explicit messaging provides the means by which typical request/response oriented functions are performed (e.g., module configuration). These messages are typically point-to-point.
I/O client	The IEEE standard for bridging 802.3 local area networks (LANs).
I/O messaging	The IEEE standard for Ethernet.
I/O server	A function that constrains flooding of multicast traffic through Layer 2 switch ports by dynamically configuring them such that multicast traffic is forwarded only to those ports associated with nodes belonging to a specific IP multicast group.
IEEE 802.1	Implicit messages are exchanged across I/O connections with an associated Connection ID (CID). The CID defines the meaning of the data and establishes the regular/repeated transport rate and transport class. No messaging protocol is contained within the message data (as with explicit messaging). Implicit messages can be point-to-point or multicast and are used to transmit application-specific I/O data. This term is used interchangeably with the term *I/O messaging*.

Term	Definition
IEEE 802.3	A function that uses the I/O messaging services of another (I/O server) device to perform a task. It initiates a request for an I/O message to the server module.
IGMP snooping	See *implicit messaging.*
implicit messaging	A function that provides I/O messaging services to another (I/O client) device. It responds to a request from the I/O client.
Internet Group Management Protocol (IGMP)	A protocol that nodes use to keep local switches apprised of their membership in multicast groups. When all nodes leave a group, switches (and routers) no longer forward packets that arrive for the group.
Layer 2 (data link layer or level)	A reference to the data link layer communication (e.g., frame formats) or data link layer connections derived from the OSI Reference Model. For local area networks, Layer 2 refers to physical frame format and addressing. Thus, for an EtherNet/IP network, a Layer 2 address is an Ethernet address.
Layer 3 (network layer or level)	A reference to the network layer communication derived from the OSI Reference Model. For TCP/IP networks, Layer 3 refers to IP and the IP packet format. Thus, a Layer 3 address is an IP address.
link	Collection of nodes with unique MAC IDs. Segments connected by repeaters make up a link; links connected by routers make up a network.
message client	A function that uses the explicit messaging services of another (Message Server) device to perform a task. It initiates an explicit message request to the server device.
message router	The object within a node that distributes messaging requests to the appropriate application objects.
message server	A function that provides explicit messaging services to another (Message Client) device. It responds to an explicit message request from the Message Client.
multicast	A transmission method by which a packet is sent to a selected subset of all possible recipients. A node can belong to one or more multicast groups. Multicast transmission is supported by Ethernet and IP protocols. Approximately half of the UDP packets on EtherNet/IP are sent via multicast.
multicast connection	A connection from one node to many. Multicast connections allow a single producer to be received by many consumer nodes.
network	A collection of one or more subnets where each node’s MAC ID is unique on the network. The series of nodes are connected by some type of communication medium. The connection paths between any pair of nodes can include repeaters, routers and gateways.
network address or node address	An identification value assigned to each node on the CIP

Term	Definition
	network. This value distinguishes a node among all other nodes on the same link. The format is network specific.
node	A collection of objects that communicate over a subnet and arbitrate using a single MAC ID. A physical device may contain one or more nodes.
node address	See *network address*.
object	(1) An abstract representation of a particular component within a product. Objects can be composed of any or all of the following components: a) data (information which changes with time); b) configuration (parameters for behavior); c) methods (things that can be done using data and configuration). (2) A collection of related data (in the form of variables) and methods (procedures) for operating on that data that have clearly defined interface and behavior.
object-specific service	A service defined by a particular object class to perform a required function that is not performed by a common service. An object specific service is unique to the object class that defines it.
octet	An octet is 8 bits that indicate no particular data type.
off-link connection	A connection between two or more devices that are on physically different networks.
originator	A node (client) that requests the creation of a connection to a target.
ping	The sequence of messages that coordinates time between consumers and producers
ping request	Request for the acknowledge of the reception of a message.
ping response	Response to a ping request.
point-to-point connection	A connection that exists between two nodes only. Connections can be either point-to-point or multicast.
port	A CIP port is the abstraction for a physical network connection to a CIP device. A CIP device has one port for each network connection. Note: network specific definitions may include additional definitions of this term within the context of the network.
produce	Act of sending data to a consumer.
producer	A node that is responsible for transmitting data.
redundant media	A system using more than one medium to help prevent communication failures.
Requested Packet Interval (RPI)	The measure of how frequently the originating application requires the transmission of data from the target application.

Term	Definition
segment	A collection of nodes connected to an uninterrupted section of physical media.
Simple Network Management Protocol (SNMP)	A standard protocol used to monitor nodes, switches, routers, and networks to which they are attached.
Spanning Tree Protocol (STP)	A switch (or bridge) protocol that uses the spanning-tree algorithm, enabling a switch to dynamically work around loops in a network topology by creating a spanning tree. Switches exchange special messages with other switches to detect loops, and then remove the loops by shutting down selected switch interfaces. Refers to both the IEEE 802.1D Spanning-Tree Protocol standard and the earlier Digital Equipment Corporation Spanning-Tree Protocol on which it is based.
subnet	A collection of nodes using a common protocol and shared media access arbitration. Subnets may have multiple physical segments and contain repeaters.
system	Contains one or more domains.
target	A destination for I/O connection or message requests. It can only respond to a request (not initiate an I/O connection or message.)
Time to Live (TTL)	A technique used in best-effort delivery systems to avoid endlessly looping packets. For example, each IP packet is assigned an integer TTL when it is created (TTL is a field in the IP packet header). Each router decrements the TTL field when the packet arrives. A router discards any packet when TTL reaches zero.
unicast	A transmission method by which a packet is sent to a single destination. All TCP packets are sent via unicast on EtherNet/IP.
Virtual LAN (VLAN)	A group of devices on one or more LANs that is configured (using management software) such that the devices can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical as opposed to physical connections, they are extremely flexible.

References

1 Comer, Douglas E., *Internetworking with TCP/IP, Volume 1: Principles, Protocol, and Architecture,* Prentice Hall, 1995.

2 *Dictionary of Internetworking Terms and Acronyms,* available on www.cisco.com

Appendix E Ingress Protection (IP) Codes

Ingress Protection (IP) Codes are environmental ratings for enclosures. The following table shows how these ratings are designated.

IP Codes

1st Digit	Protection Against Foreign Objects	2nd Digit	Protection Against Moisture
0	Not protected	0	Not protected
1	Protected against objects greater than 50 mm	1	Protected against dripping water
2	Protected against objects greater than 12 mm	2	Protected against dripping water when tilted up to 15 N
3	Protected against objects greater than 2.5 mm	3	Protected against spraying water
4	Protected against objects greater than 1.0 mm	4	Protected against splashing water
5	Dust protected	5	Protected against water jets
6	Dust tight	6	Protected against heavy seas
	------	7	Protection against the effects of immersion
	------	8	Protection against submersion

Example marking: IP 68 would indicate a dust-tight (first digit 6) piece of equipment that is protected against submersion in water (second digit 8).

Source: Underwriters Laboratories, Inc.