

Industrial Ethernet: A Control Engineer's Guide

Abstract

As part of a continuing effort to make their organizations more efficient and flexible, manufacturers are rapidly migrating to Industrial Ethernet technology to network their industrial automation and control systems and connect plant and business systems to achieve business objectives. Industrial Ethernet enables a more flexible, responsive system that encompasses both real-time data from the production floor. This end-to-end networking architecture provides connectivity, collaboration, and integration from the device level to enterprise business systems.

This white paper provides an overview of Ethernet technology and its benefits in the data networking environment. It discusses the requirements and benefits of using a switched Ethernet architecture in industrial networking environments. Benefits include:

- Real-time network performance (including low latency, low jitter, and minimal packet loss) to develop deterministic systems
- Security
- Reliability
- Manageability and ease-of-use features
- Ability to add innovative technologies such as voice, video and collaboration

Industrial Ethernet applies the Ethernet and IP suite of standards developed for data communication to manufacturing control networks. By implementing an intelligent Industrial Ethernet solution, organizations can build a manufacturing infrastructure that delivers the resiliency and network security of traditional fieldbus solutions, as well as the improved bandwidth, open connectivity, and standardization that Ethernet provides. Industrial Ethernet gives organizations substantially greater control over their networked manufacturing equipment.

Introduction

Manufacturing companies are increasingly expanding their global operations to address new opportunities and reduce operating costs. They are also seeking to continuously improve efficiency and drive down costs for existing facilities and processes. Achieving the goals of globalization and operations excellence requires improving connectivity between plant and business systems for real-time visibility to information and effective collaboration. This helps to assure consistent quality and performance across global operations, and to reduce the cost of design, deployment, and support of distributed manufacturing and IT systems. Manufacturers must be able to balance production with demand to optimize material usage and asset utilization, while continuing to meet increasingly exacting customer requirements and metrics for on-time delivery. Manufacturers also need to improve response to events that occur on the plant floor, regardless of location, while implementing more flexible and agile operations in order to react to rapidly changing market conditions.

As manufacturers seek to improve processes, increase productivity, reduce operating costs, and integrate manufacturing and business networks, many are turning to Ethernet technology on the factory floor. This migration is rapidly gaining momentum. According to a recent ARC Advisory Group study¹, the worldwide market for Industrial Ethernet devices is expected to grow at a rate of approximately 30 percent compound annual growth rate (CAGR)

¹ Industrial Ethernet Infrastructure Worldwide Outlook, Market Analysis and Forecast through 2011; ARC Advisory Group, September 2007

through 2011. Once considered a solution that was limited to corporate network environments, Ethernet technology has proven to be a robust alternative that can meet the unique needs of the manufacturing environment.

Industrial Ethernet networks that use intelligent switching technology can offer a variety of advantages compared to traditional industrial networks. The technology can be deployed using a switched Ethernet architecture and has proven successful in multiple critical applications in different markets. Because the technology is based on industry standards, Industrial Ethernet enables organizations to save money by moving away from expensive, closed, factory-floor-optimized networks. Using standard Ethernet technologies also reduces overall risk and provides investment protection, as manufacturers and automation vendors can take advantage of continued industry investment and innovation in compatible technologies (such as 1 GB, 10 GB Ethernet, and video and voice over IP). By providing a scalable platform that can accommodate multiple applications, Ethernet-based automation systems can increase flexibility and accelerate deployment of new applications in the future. At the same time, Ethernet delivers the network security, performance, and availability required to support critical manufacturing applications.

To deploy this technology, engineers on the manufacturing floor should be familiar with some of the important concepts behind Industrial Ethernet. This paper will provide a general overview of the most important traditional Ethernet technologies in use today. It will also discuss how Industrial Ethernet upgrades traditional, proprietary factory-floor networks to a low-cost, high-performance, scalable architecture. Finally, this paper will review some of the intelligent features that make Industrial Ethernet an attractive choice for manufacturing organizations.


Traditionally Separate Networks

Today, many manufacturing companies maintain separate networks to support their factory floor operations and business operations (Figure 1). Over the years, these networks were developed to respond to the different information flows and control requirements involved with manufacturing processes.

The corporate IT network supports traditional administrative functions and corporate applications, such as human resources, accounting, and procurement, as well as WAN connections between sites and Internet connectivity. This network is usually based on the Ethernet and IP suite of standards.

The control-level network connects control and monitoring devices, including programmable automation controllers, PC-based controllers, I/O racks, drives, and human-machine interfaces (HMI). This network, which typically has not been based upon standard Ethernet and IP in the past, requires a router or, in most cases, a gateway to translate application-specific protocols to Ethernet-based protocols. This translation lets information pass between the control network on the factory floor and the corporate network infrastructure, but has limited functionality and bandwidth, and requires significant effort to keep up to date.

The device-level network links the controllers with the plant floor's I/O devices, including sensors such as transducers, photoeyes, and flowmeters, and other automation and motion equipment, such as robotics, variable frequency drives, and actuators. Interconnectivity between these devices was traditionally achieved with a variety of fieldbuses such as DeviceNet, Profibus, and Modbus. Each fieldbus has specific power, cable, and communication requirements, depending on the factory application it supports. This has led to a replication of multiple networks in the same space and the need to have multiple sets of spares, skills, and support programs within the same organization.

Figure 1. Traditional Proprietary Fieldbus Architecture

Instead of using architectures composed of multiple separate networks, Industrial Ethernet can unite a company's administrative, control-level, and device-level networks to run over a single network infrastructure. In an Industrial Ethernet network, fieldbus-specific information that is used to control I/O devices and other manufacturing components are embedded into Ethernet frames. Because the technology is based on industry standards rather than on custom or proprietary standards, it is more interoperable with other network equipment and networks.

Increasing Industry Support

Industrial Ethernet technology is rapidly being embraced by multiple organizations and vendors, including the Industrial Ethernet Association (IEA), the Open DeviceNet Vendor Association (ODVA), Modbus.org, Fieldbus Foundation, and Profinet and Profibus International (PI).

Adoption of Industrial Ethernet by end users has also been significantly increasing over the last few years. ARC Research predicts approximately 30 percent CAGR through 2011 for Industrial Ethernet nodes shipped, and in 2006, ODVA announced² that over 1 million Ethernet/IP nodes had been shipped to date.

The OSI Reference Model


At the heart of data networking is the Open Systems Interconnection (OSI) reference model. This conceptual model describes how information from a software application in one computer moves through a network medium to a software application in another computer. The model was developed by the International Organization for Standardization (ISO) in 1984, and it is now considered the primary architectural model for intercomputer communications.

² ODVA press release April 24, 2006:

<http://www.odva.org/Home/ABOUTODVA/PressRoom/NewsArchive/tabid/123/ctl/Details/mid/507/ItemID/3/Default.aspx>

The OSI reference model divides the tasks involved in moving information between networked computers into seven smaller, more manageable task groups. These tasks are then assigned to seven layers in the OSI model. Each layer is self-contained so that the tasks assigned to it can be implemented independently. Figure 1 shows the seven OSI layers.

Figure 2. Layers of the OSI Reference Model



Functions of the OSI Layers

The seven layers of the OSI reference model can be divided into lower layers (1–4) and upper layers (5–7). The lower layers focus on data-transport functions while the upper layers focus on the applications. 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, such as network cabling. Ethernet resides in Layer 2, as do some implementations of traditional fieldbuses such as DeviceNet, which uses the Controller Area Network (CAN) protocol. Layer 3 takes care of the logical addressing and routing (which way to send data). Its most common implementation uses the Internet Protocol (IP), which is the core of World Wide Web addressing and routing. Layer 4, the last of the lower layers, is the transport layer. It ensures that data is delivered error-free and in the correct sequence. Industrial Ethernet is broader than traditional Ethernet technology. While Ethernet technology refers only to Layers 1 and 2, most Industrial Ethernet solutions also encompass Layers 3 and 4, using IP addressing in Layer 3, and Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) in Layer 4, in what is referred to as the IP suite.

The upper layers of the OSI reference model are responsible for application tasks and are usually 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 involve network communications.³

³ For more information about the OSI reference model, visit:
<http://www.cisco.com/en/US/docs/internetworking/technology/handbook/Intro-to-Internet.html>

What Is Ethernet?

Ethernet is by far the most widely used LAN technology today, connecting more than 85 percent of the world's LAN-connected PCs and workstations. Ethernet refers to the family of computer networking technologies covered by the IEEE 802.3 standard, and can run over both optical fiber and twisted-pair cables. Over the years, Ethernet has steadily evolved to provide additional performance and network intelligence. This continual improvement has made Ethernet an excellent solution for industrial applications. Today, the technology can provide four data rates.

- **10BASE-T Ethernet** delivers performance of up to 10 Mbps over twisted-pair copper cable.
- Fast Ethernet delivers a speed increase of 10 times the 10BASE-T Ethernet specification (100 Mbps) while retaining many of Ethernet's technical specifications. These similarities enable organizations to use 10BASE-T applications and network management tools on Fast Ethernet networks.
- **Gigabit Ethernet** extends the Ethernet protocol even further, increasing speed tenfold over Fast Ethernet to 1000 Mbps, or 1 Gbps. Because it is based upon the current Ethernet standard and compatible with the installed base of Ethernet and Fast Ethernet switches and routers, network managers can support Gigabit Ethernet without needing to retrain or learn a new technology.
- **10 Gigabit Ethernet**, ratified as a standard in June 2002, is an even faster version of Ethernet. Because 10 Gigabit Ethernet is a type of Ethernet, it can support intelligent Ethernet-based network services, interoperate with existing architectures, and minimize users' learning curves. Its high data rate of 10 Gbps makes it a good solution to deliver high bandwidth in WANs and metropolitan-area networks (MANs).

More than 300 million switched Ethernet ports have been installed worldwide. Ethernet technology enjoys such wide acceptance because it is easy to understand, deploy, manage, and maintain. Ethernet is low-cost and flexible, and supports a variety of network topologies. Compared to traditional, non-Ethernet-based industrial solutions that have a data rate of between 500 Kbps to 12 Mbps, Ethernet technology can deliver substantially higher performance. And, because it is based on industry standards, it can run and be connected over any Ethernet-compliant device from any vendor.

What Are IP, TCP, and UDP?

IP is often understood as the Internet Protocol suite—a suite of protocols and standards on which the Internet and most enterprise networks are based. The IP suite includes not only lower-level specifications, such as TCP and IP, but specifications for such common applications as e-mail, terminal emulation, and file transfer. The most relevant of these to the factory floor though are IP itself, TCP, and UDP.

IP is the common and primary network (Layer 3) protocol in the Internet suite. IP represents the core of the Internet Protocol suite. It defines an address by which the network can transmit the packet from source to destination—even across LANs. This is opposed to the MAC addresses, which are used to network locally, within a LAN. IP has two primary responsibilities: providing connectionless, best-effort delivery of datagrams or packets through a network; and providing fragmentation and reassembly of datagrams to support data links with different maximum-transmission unit (MTU) sizes. More on IP address allocation and administration is found in the network management section.

TCP provides reliable, in-order delivery of packets between two devices. It relies upon IP. TCP applications establish connections between one another over which they send packets. TCP is a stateful protocol; it maintains the state after the packet is sent. TCP checks whether all packets have arrived and can request re-transmission if a packet is dropped, lost, or corrupted during transmission. Due to this overhead though, TCP is not always ideal for real-time applications.

UDP is often used for real-time communications such as voice and I/O traffic. UDP also relies on IP. UDP does not guarantee delivery or the order of the packets, thus simplifying the protocol. The applications would rather drop a packet than receive it late. UDP is considered a “stateless” protocol. It is compatible with [packet broadcast](#) (sending to all on local network) and [multicasting](#) (sending to all subscribers).


There are a whole host of other protocols designed for specific purposes in the networking world, but the key ones for an automation and control system are IP, TCP, and UDP. For more information on the OSI model and networking protocols, refer to the Internetworking Technology Handbook from Cisco:

http://www.cisco.com/en/US/tech/tk1330/tsd_technology_support_technical_reference_book09186a00807594e5.html

What Is Industrial Ethernet?

Recognizing that Ethernet is the leading networking solution, many industry organizations are porting the traditional fieldbus architectures to Industrial Ethernet. Industrial Ethernet applies the Ethernet standards developed for data communication to manufacturing control networks (Figure 3). Using IEEE standards-based equipment, organizations can migrate all or part of their factory operations to an Ethernet environment at the pace they wish.

Figure 3. Using Intelligent Ethernet for Automation Control



For example, Common Industrial Protocol (CIP) has implementations based upon Ethernet and the IP protocol suite (EtherNet/IP), DeviceNet, and ControlNet (among others). Most controllers (with appropriate network connections) can transfer data from one network type to the other, leveraging existing installations, yet taking advantage of Ethernet. The fieldbus data structure is applied to Layers 5, 6, and 7 of the OSI reference model over Ethernet, IP, and TCP/UDP in the transport layer (Layer 4).

The advantage of Industrial Ethernet is that organizations and devices can continue using their traditional tools and applications running over a much more efficient networking infrastructure.

Industrial Ethernet not only gives manufacturing devices a much faster way to communicate, but also gives the users better connectivity and transparency, enabling users to connect to the devices they want without requiring separate gateways.

Technology Tailored for Manufacturing

Although Industrial Ethernet is based on the same industry standards as traditional Ethernet technology, the implementation of the two solutions is not always identical. Industrial Ethernet usually requires equipment that can handle more severe environmental conditions, flexible node counts, varieties of media, very predictable real-time data traffic performance, and increased levels of segmentation as compared to traditional Ethernet networks in a corporate data network.

The primary difference between Industrial Ethernet and traditional Ethernet is the type of hardware used. Industrial Ethernet equipment is designed to operate in harsh environments. It includes industrial-grade components, convection cooling, and relay output signaling. And it is designed to operate at extreme temperatures and under extreme vibration and shock (and other conditions). Power requirements for industrial environments differ from data networks, so the equipment runs using 24 volts of DC power. To maximize network availability, Industrial Ethernet equipment also includes fault-tolerant features such as redundant power supplies. The equipment is also modular in order to meet the highly varying requirements of a factory floor.

Other differences are important as well. The Industrial Ethernet automation and control protocols themselves and their use of the available technologies within the standard Ethernet and IP protocol suite often differ significantly from standard Ethernet implementations. For example, in most automation and control applications, 80 percent of the network traffic is local—one local device talking to another local device often using multicast (one sender, many receivers) packets. In most IT installations, the reverse is true where 80 percent of the network traffic is routed to external locations (such as the data center or the Internet) using unicast (one sender, one receiver) packets. Automation and control systems also differ from other applications in their need for determinism and real-time network requirements—quick and consistent transmission of the data. The Industrial Ethernet network must be designed and implemented with these differences in mind.

Ethernet and the IP protocol suite have developed a number of technologies and features that support these requirements. To help optimize synchronous data access, Industrial Ethernet equipment must include the intelligence to support features such as multicast control (IGMP Snooping), QoS, and virtual LANs (VLANs). Other high availability, security, and management functions should also be considered depending on the specific automation and control application.

Benefits of a Switched Ethernet Architecture

Organizations can choose from a variety of devices and architectures when building an Ethernet LAN. Devices range from simple hubs, to unmanaged switches, to intelligent, managed switches. The network components are important to the proper functioning of the automation network, and careful consideration should be given to selecting the appropriate device.

Early Ethernet deployments often used hubs. Hubs act at the physical layer of the OSI model, and are essentially repeaters that connect multiple devices over the same shared medium. Because hubs use a shared medium, collisions can occur when multiple devices try to communicate at the same time. This may not be a significant concern in a small network without high-performance requirements, but is typically not acceptable in environments where real-time, predictable performance is important, such as automation networks. The collisions that resulted from the use of hubs contributed to the perception that Ethernet is not deterministic, even though hubs are rarely used anymore.

Over the past 10 years, most Ethernet deployments have used full-duplex switched Ethernet switches. Switches 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, there are no hubs so each Ethernet network has a dedicated segment for every node. Because 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 data over to the appropriate segment. In a fully switched network, nodes only communicate with the switch and never directly with each other.

Fully switched networks employ either twisted pair or fiber-optic cabling, both of which use separate conductors for sending and receiving data. The use of dedicated communication channels allows nodes to transmit to the switch at the same time the switch transmits to them, eliminating the possibility of collisions. Transmitting in both directions also can effectively double the apparent speed of the network when two nodes are exchanging information. For example, if the speed of the network is 10 Mbps, each node can transmit at 10 Mbps at the same time.

Switches usually work at Layer 2 (data link) of the OSI reference model using MAC addresses, and deliver a number of important advantages compared to hubs and other LAN devices. Some of these advantages include the following:

- **Predictable performance:** The ability to ensure that a packet is sent and received in a specific period of time is an important design goal for industrial networks. For the network to support predictable, real-time traffic, the design must be as simple and highly structured as possible.
- **Latency:** Switches normally have very low latencies, which refers to the time it takes for a network packet to transit between a source and a target. Most control operations in industrial applications can tolerate latencies of 10 to 50 milliseconds (ms). Because control traffic frames in industrial applications are usually below 500 bytes, the latency introduced by a switch at 100 Mbps is only about 30 microseconds with a worst-case scenario of close to 100 microseconds—well below the limit and 100 times faster than most applications require.
- **Standardization:** One of the main motives for using Industrial Ethernet is the need to standardize around a common infrastructure. Unlike proprietary technologies that often tie companies to a particular vendor, standardized solutions free users to choose the best application for a given solution. And a standard Ethernet network brings to the factory floor the economies of scale enjoyed by today's large base of Ethernet users, lowering costs and increasing the number of potential equipment vendors and products.

Managed switches provide performance, management, diagnostics, and security capabilities that are not supported on unmanaged switches. These types of features allow the network administrator to configure the switch to provide traffic prioritization, basic and advanced security capabilities, multicast traffic control, diagnostic capabilities, and a number of other features that are important for most industrial and office network environments. Given the critical nature and performance requirements of automation and control networks, a managed switched Ethernet architecture is the most appropriate choice for most industrial environments. Some of the most important features on intelligent managed switches in an industrial environment include:

- **Packet loss under congestion:** Today's intelligent switches offer QoS features that make it possible to prioritize critical traffic so that it will be handled with priority and not be dropped due to congestion. By implementing simple QoS parameters in an intelligent switch, organizations can prioritize critical traffic over non-critical traffic at wire speed, helping to ensure consistent packet delivery and integrity for the control network. Even under heavy congestion, QoS features help ensure that important traffic will reach its destination quickly and consistently.
- **Broadcasts and multicast:** Industrial applications often rely on broadcast or multicast communication. Intelligent switching platforms can dynamically configure the interfaces so that traffic is forwarded only to ports associated with requested data. This feature reduces the load of traffic crossing the network and relieves the client devices from processing unneeded frames.
- **Network analyzers:** Intelligent switches allow traffic analyzers to remotely monitor any port in a network (also known as port mirroring), which saves organizations time and money and reduces the amount of hardware that must be deployed to monitor and optimize network usage.

- **Security:** Managed switches play a major role in a security approach as the first point of access to a network and system. It starts with port security and settings that control which devices can connect. And managed switches can be configured to reduce or eliminate common types of attacks (intentional or unintentional) such as broadcast or multicast storms. A broadcast or multicast storm results when a device produces a large quantity of broadcasts or multicast messages that flood the network. With a managed switch you can apply VLANs and access control lists (ACLs) to segregate devices and traffic from one another, even down to a port level.
- **Diagnostics:** A critical factor when resolving a problem on the factory floor is having the right information. Managed switches provide a host of diagnostic information that can be helpful to resolve network and device issues occurring in the automation and control systems. Critical diagnostic information includes port status, amount of traffic being passed, and the ability to mirror a port to see the type of traffic a node is generating or receiving. As well, switches are being developed (and are already available) that function as common industrial devices that can be directly managed and controlled by the automation and control applications, like any other device on the factory floor. This enables production personnel to monitor and maintain the network infrastructure as they do with the other automation and control equipment.

Network Requirements: The Need for Intelligence

When implementing an Industrial Ethernet solution, companies should be careful to select Ethernet products that offer the intelligent features required to support manufacturing applications. Network intelligence enables organizations to build a manufacturing infrastructure that matches the resiliency and network security of traditional fieldbus solutions, while at the same time providing the benefits of higher bandwidth, open connectivity, and standardization offered by Ethernet-based platforms. The important qualities behind an intelligent Industrial Ethernet solution include network security, reliability, determinism, and manageability (ease of use).

Network Security

Ethernet technology can provide not only excellent performance for manufacturing applications, but a wide range of network security measures to maintain availability, integrity, and confidentiality of the automation and control systems. Availability is most often cited as the key security requirement from a manufacturing point of view: keep the automation and control systems operational. Integrity protects data and systems from intentional or accidental alteration. Confidentiality helps ensure that data cannot be accessed by unauthorized users. These network security advantages protect manufacturing devices such as programmable automation controllers (PACs) as well as PCs, and apply to both equipment and data security.

As with any system characteristic, security is maintained through a lifecycle of design, implementation, maintenance, and improvement. Security and administration policies are a key foundation for developing robust network security. A security policy should logically segment the devices and network in a manufacturing environment into groups or zones, on which the policies can be applied. Once the security policy is defined, there are a number of key technical capabilities available to implement the policy. These include, but are not limited to:

- **VLAN configuration:** A VLAN is a group of devices on one or more physical LANs that are configured by the network so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Conversely, VLANs enable devices on a single LAN segment to be logically segmented into different VLANs. VLANs should be used to segment devices that need to communicate to each other. Then, other devices or users that infrequently need to communicate with those devices can be allowed access to the VLAN. VLANs form a basic level of segmentation on which a security policy can be applied.

- **Access control and authentication:** Access control is commonly implemented using firewalls or network-based controls protecting access to critical applications, devices, and data so that only legitimate users and information can pass through the network. However, access-control technology is not limited to dedicated firewall devices. Any device that can make decisions to permit or deny network traffic, such as an intelligent switch, is part of an integrated access-control solution.
- **Firewalls:** A firewall regulates network traffic between various networks. Except for completely disconnecting networks, firewalls are about the strongest form of segmentation. Firewalls inspect all aspects of traffic flowing between networks, even inspecting the data content of a packet (versus just the header information)—a process known as deep packet inspection—and maintaining stateful information of the network traffic. A firewall can be hardware- or software-based. Firewalls are applied at major risk points in networks, for example where manufacturing networks interface with enterprise networks or between the enterprise network and the Internet.
- **Demilitarized zone (DMZ):** A DMZ is a buffer zone between various areas of a network, and supports sharing of data and services between the network zones. A common example is the DMZ between an enterprise's Internet presence and the enterprise network. A DMZ supports exchange of data and services between the zones, yet enables strict control of the traffic from either zone. Cisco recommends that a manufacturing network supporting automation and control systems be segmented from the general enterprise network with a DMZ.
- **Secure connectivity and management:** To provide additional protection for manufacturing networks, organizations can take several approaches to authenticate and encrypt network traffic. Using VPN technology, Secure Sockets Layer (SSL) encryption can be applied to application-layer data in an IP network. Organizations can also use IP Security (IPsec) technology to encrypt and authenticate network packets to thwart network attacks such as sniffing and spoofing.

Manufacturers can use all or a portion of these technologies to help ensure network availability, confidentiality, and integrity.

For a more in-depth look at network security in the manufacturing space, please refer to the Cisco white paper, "Ethernet to the Factory: Securing Today's Global Networks in Industrial Environments"⁴.

Reliability

Because factory-floor applications run in real time, the network must be available to users on a continuous basis, with little or no downtime. Manufacturers can help ensure network reliability using effective network design principles, as well as intelligent networking services.

Reliability needs to be considered at each layer in the OSI model. Especially important for network design considerations are Layers 1–3. At each layer, a number of techniques or considerations can be applied to achieve the needed level of reliability (Table 1).

Table 1. Network Design Considerations

OSI Layer	Considerations
Physical	• Redundant components (power supplies, supervisors, etc.) • Redundant devices (switches and routers)
Data Link	Network topology: • Redundant links • Redundant paths
Network	• IP routing

⁴ Available on the Ethernet to the Factory home page: http://www.cisco.com/web/strategy/manufacturing/ettf_overview.html

Reliability exists at other layers, but is applied by the automation and control applications. For example, TCP (a Layer 4 protocol) is inherently reliable because dropped or corrupt packets are automatically resent when detected by the protocol. But the overhead and delay represented in that process have led automation and control vendors to use UDP for some types of traffic. UDP has less overhead, but no automated resending for particular information to communicate and therefore relies upon the application to detect and manage packet loss or delay.

Physical Layer Reliability


At the physical layer, a number of techniques can be applied to help achieve a resilient, highly available network. First, the various components can be configured or purchased with resilient features such as redundant power supplies (or even UPS), and redundant components (such as fans, CPUs, network interface cards [NICs], etc.). Additionally, some devices may also support in-line upgradeability of components or software that allows for continued service while the device is being maintained or upgraded. These techniques will usually significantly improve the mean time to repair (MTTR) the device itself or ensure the device has network access in the case of a media disruption or port failure (on either the end device or switch).

Using redundant devices may also help maintain high network availability. For example, multiple switches or routers can be configured in a high-availability manner so that in the case of disruption of one device, the other device will take over the network services quickly and automatically. There is a wide variety of solutions, both standard and proprietary, that may or may not allow for hot (working) or standby (on but not performing) operations of both devices and re-establish network communications faster or slower (also known as network convergence).


Data Link Reliability

Manufacturers deploying an Ethernet solution should design networks with redundant paths to ensure that a single device or link outage does not take down the entire network. How the end devices and network devices are all interconnected is a network topology. Two network topologies most often used to achieve higher availability are ring and redundant star. The topology chosen also has implications on wiring cost and complexity, performance, and installation and maintenance cost. Other topologies (such as bus or trunk-drop) may be cheaper to install and easier to maintain, but are more susceptible to outage and have a higher impact when a connection or device is lost.

In redundant star designs (Figure 4), switches and routers are connected in a hierarchical fashion. The first layer where devices are connected to switches is often referred to as the access layer. These switches provide connections for endpoint devices such as PLCs, robots, and HMIs. Access-layer switches generally operate at Layer 2 (data link) of the OSI model. Above the access layer is another layer of switches referred to as the distribution layer. These switches interlink the various access layer switches. If they support multiple cell/area zones, they may need to operate at Layer 3 (network) of the OSI model, referred to as Layer 3 switches or routers, to support multiple VLANs.

Figure 4. Redundant Star Network Topology

In ring topologies (Figure 5), all devices are connected in a ring. Each device has a neighbor to its left and right. If a connection on one side of the device is broken, network connectivity can still be maintained over the ring via the opposite side of the device. In a typical topology, the ring is at the access layer, and connected up through the distribution and core layers using a redundant star topology. In this model, the distribution and core layers provide the same functionality as in the redundant star, with the distribution layer routing between cells, and the core connecting to higher-level or external networks.

Figure 5. Ring Topology

It is important to understand the physical environment and network performance requirements when determining the optimal network topology. The physical layout of automation networks is often different than a traditional IT environment, which may drive the use of the ring topology. Unlike traditional IT networks, which are largely redundant star topology networks, manufacturing networks sometimes have physical limitations, based on factory layout and equipment design, which drives the use of the ring topology, due to the cost and complexity of cabling a dual redundant star down a long production line.

While the ring topology does help address some of the challenges with cabling complexity in manufacturing environments, it is not without network performance tradeoffs. A redundant star offers the following advantages:

- **Resiliency:** A redundant star can recover from multiple link failures.
- **Convergence:** The network will reconverge more quickly in a redundant star configuration, minimizing or eliminating the interruption of traffic flow when a link fails.
- **Consistent performance:** A consistent number of hops (and typically fewer hops than a ring topology) provides more predictable performance and real-time characteristics that are often important in an automation and control network.
- **Fewer bottlenecks:** Reduces the chances of segment oversubscription. While many manufacturing applications don't require significant bandwidth, as more applications or technologies are added over time, network performance could degrade due to network bottlenecks.

All of these network performance implications need to be considered along with the cost, complexity, and feasibility of cabling when determining the optimal network topology for a specific application in a manufacturing environment.

Table 2 highlights the appropriate topology based on some common concerns:

Table 2. Recommended Topologies

Key Concerns	Topology to Use
• Highly available network with minimal convergence • High-performance network with minimal bottlenecks • Straightforward network design	Redundant star
• Cabling complexity is a major concern • Highly available network is important • Cost is important	Ring
• Cost and simplicity over availability and performance	Trunk-drop or bus

Spanning Tree Protocols

When a network has redundant paths, a loop can be formed in the network if the appropriate protocols are not used to block redundant paths until a failure is detected. In Ethernet networks, an unmanaged loop is dangerous because broadcast and multicast messages are continuously passed until the network overloads, called a broadcast storm. A variety of protocols exist to prevent loops from being formed in the network when devices are interconnected via multiple paths by logically closing a connection or path until a failure is detected and the connection or path is re-enabled. The standard protocol for this function is the Spanning Tree Protocol and its more modern version called Rapid Spanning Tree Protocol (RSTP). Spanning Tree Protocol was developed to manage redundant paths in a Layer 2 network. These protocols virtually close a link or links to eliminate loops, yet maintain network viability. If a problem occurs on a network node, the protocol detects it and re-opens any needed closed links to re-establish network connectivity. This function is referred to as convergence. The protocols also automatically detect the repair of broken links and subsequently re-converge the network.

The original Spanning Tree Protocol has been considered too slow for industrial environments. To address these performance concerns, the IEEE standards committee has ratified the Rapid Spanning Tree Protocol (802.1w). This

protocol provides sub-second convergence times that vary between 500 and 2000 ms, depending on network topology and size. Using 802.1w, organizations can achieve the benefits of redundant Ethernet networks, with the performance and reliability that many manufacturing applications demand.

Other Resiliency Options

Some industrial automation and control applications require faster convergence times than can be achieved with Rapid Spanning Tree Protocol, or different variations and extensions of RSTP, including rapid PVST+ and MSTP. Applications that require this rapid convergence time may require different convergence protocols. There are a number of different protocols that are appropriate for different applications and topologies. A few of these are detailed below:

- **EtherChannel:** EtherChannel is a link aggregation protocol that is designed to provide additional bandwidth between two devices by aggregating multiple Ethernet connections. Due to the fast recovery from a link failure, this protocol can be used to achieve resiliency with fast recovery between two different devices in an industrial network. EtherChannel has the advantages of being standards based, and can achieve convergence times of less than 100ms. Due to the nature of this protocol, it is limited to topologies with redundant links, such as a redundant star, and cannot be used in topologies with only one link between switches, such as a ring.
- **Flex Links:** Flex Links defines an active and standby link between two networking devices (access and distribution switch), and switches all traffic to the standby port if a failure occurs on the active link. Like EtherChannel, Flex Links is used in a redundant star topology, and can achieve convergence times of less than 100ms. Flex Links is a Cisco® proprietary protocol, but can interoperate and coexist with other protocols such as STP.
- **Resilient Ethernet Protocol (REP):** Resilient Ethernet Protocol (REP) is a Cisco technology that is designed to achieve fast convergence in a ring topology or a variety of complex topologies. REP can achieve convergence times of less than 50ms in a ring topology of up to 50 switches. REP is Cisco proprietary, but is a segment technology and can interoperate with other protocols such as STP.

The appropriate convergence protocol depends on the environment. The topology, application requirements, network management considerations, and need to support or interoperate with open standards will determine the best option for a specific industrial automation and control system network.

Network (Layer 3) Reliability

Another form of network reliability exists in the IP or network layer. Although much of the Industrial Ethernet traffic is local and never really relies on Layer 3 networking, this reliability is important for the overall network availability. Every device in an Industrial Ethernet network has an IP address. Essentially a Layer 3 switch or router is required to route traffic based upon IP addresses. This functionality is needed to route traffic between VLANs—essentially based upon a packet's destination IP address.

There is a number of standard and proprietary protocols that have been developed that allow Layer 3 switches and routers to reliably route packets to their destination based upon IP address. These protocols allow various devices to communicate and maintain viable routes between each other so that packets can always be forwarded if a viable path exists even as connections or devices fail. Typically, the choice of routing protocol and the configuration and maintenance of the Layer 3 switches and routers are performed by networking experts in IT organizations. For more information on routing protocols and best practices, please refer to the Cisco's Converged Plantwide Ethernet.⁵ This

⁵ View the Converged Plantwide Ethernet Design and Implementation Guide at:
http://www.cisco.com/en/US/docs/solutions/Verticals/CPwE/CPwE_DIG.html

guide was developed with Rockwell Automation and provides validated design guidance and best practices for factory automation networks based on Ethernet.

Real-Time Traffic Performance

Because manufacturing processes depend on the precise synchronization of processes, the network must be optimized to deliver consistent, real-time performance. To achieve the necessary performance, technologies that prioritize and filter traffic, and that segment the network, need to be part of the network design. Data must be prioritized using QoS to ensure that critical information is received first. Multicast applications that are prevalent in manufacturing environments must be well managed using Internet Group Management Protocol (IGMP). IGMP manages multicast traffic by establishing a “publish and subscribe” mechanism. With IGMP, switches can process or “snoop” multicast traffic and determine which devices have subscribed to which multicast groups and send the traffic only to those devices that want the packets. And the devices and controllers must be grouped appropriately to optimize the flow of traffic within and between different cells.

The Producer-Consumer Model in Industrial Ethernet

Many Industrial Ethernet applications depend on IP multicast technology. IP multicast allows a host, or source, to send packets to another group of hosts called receivers anywhere within the IP network using a special form of IP address called the IP multicast group address.

Many Industrial Ethernet environments use a producer-consumer model, where devices generate data called “tags” for consumption by other devices. The devices that generate the data are producers and the devices receiving the information are consumers. In this scenario, multicast is more efficient than unicast, because consumers will often want the same information from a particular producer, yet the producer only has to send the information once. Each device on the network can be both a producer and a consumer of traffic. Although this is a significantly different application of multicast than IT-based multicast applications (such as video or multimedia), the same standard mechanisms developed to optimize this traffic apply very well.


While most devices generate very little data, networks with a large number of nodes can generate a large amount of multicast traffic, which can overrun end devices in the network if left unmanaged. By using switching infrastructure that supports QoS, Internet Group Management Protocol (IGMP) snooping, and VLANs, organizations can control and effectively manage multicast traffic in manufacturing environments.

IGMP Snooping

Many manufacturing applications depend on multicast traffic, which can introduce performance problems in the network and in the end devices themselves. To address these challenges in an Industrial Ethernet environment, organizations can turn on IGMP “snooping”⁶ on their managed switch network. IGMP snooping limits the flooding of multicast traffic by dynamically configuring the interfaces so that multicast traffic is forwarded only to 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 the traffic. This is very important because it reduces the load of traffic traversing through the network. It also relieves the end devices from processing frames that are not needed.

In a producer-consumer model used by Industrial Ethernet protocols such as CIP, IGMP snooping can limit unnecessary traffic from the I/O device that is producing, so the traffic only reaches the device consuming that data. Messages delivered to a particular device that were intended for other devices consume resources and slow performance, so networks with many multicasting devices will suffer performance issues if IGMP snooping or other multicast limiting schemes are not implemented.

⁶ The Internet Group Management Protocol (IGMP) is a communications protocol used to manage the membership of IP multicast groups.

Figure 6. The Importance of IGMP Snooping in Producer-Consumer Model

The IGMP snooping feature allows Ethernet switches to “listen” to the IGMP conversation between end devices. With IGMP snooping, the Ethernet switch examines the IGMP traffic coming to the 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 the switch relays the IGMP queries, it deletes entries periodically if it does not receive any IGMP membership reports from the multicast clients. Only devices that are part of the group receive the multicast messages sent to that group, thus reducing the amount of messages that the network must send and that the end devices need to process.

Quality of Service


An Industrial Ethernet network may transmit many different types of traffic, from routine data to critical control information (such as I/O traffic), or even bandwidth-intensive video or voice. The network must be able to distinguish among and give priority to different types of traffic. By giving priority to different types of traffic, the network can deliver real-time network services: low latency and jitter and minimal packet loss when the network infrastructure is under load. This capability to share the network with other applications, yet maintain the priority of the critical traffic, is a key differentiating factor for Industrial Ethernet versus existing industrial network protocols.

Organizations can implement QoS using several techniques. QoS involves three important steps. First, different traffic types in the network need to be identified through classification techniques and then tagged. The classification and marking can occur either in the end device or the network infrastructure, depending on the capabilities of both. There are also various ways to tag and priority levels that can be applied. It is important to understand the type of traffic that will exist on the network and ensure that the QoS approach takes into consideration all of those traffic types and the capability of the network infrastructure. Second, advanced buffer-management techniques need to be implemented to prevent high-priority traffic from being dropped during congestion. Finally, scheduling techniques need to be incorporated to transmit high-priority traffic from queues as quickly as possible. All three steps need to be considered when developing a technique to implement QoS.

Although the application of QoS should be designed and tested before implementation, it is relatively easy to deploy and maintain through the use of predefined port and switch configurations. As well, the QoS approach for the manufacturing zone may be different for the approach in the enterprise zone (where, for example, voice traffic gets

the highest priority), thus reinforce the need to logically segment the manufacturing zone to set up boundaries between the different QoS implementations.

Figure 7. Applying QoS to Industrial Applications



In an Industrial Ethernet application, real-time I/O control traffic would share network resources with administrative data (such as explicit data), configuration files (FTP, for example) and data-collection flows, as well as other traffic, in the upper layers of the OSI reference model. By using QoS to give high priority to real-time UDP control traffic, organizations can realize the benefit of sharing resources yet maintain the real-time network characteristics required for I/O control traffic in industrial automation and control applications.

Network Management and IP Addressing

Network Management

Network management is a set of tools, applications, and devices used to monitor and maintain a network. Although a typical automation and control network does not drastically change after deployment, as with all other aspects of the system, the network does need to be maintained and managed. To a large extent, these functions have not been incorporated into the automation and control systems, but this is changing. Therefore, production control engineers designing and managing automation and control systems need to be aware of how the network will be managed and what tools, training, and resources are required to put those in place.

The key functions of a network management as defined by the International Organization for Standards (ISO) are:

- **Performance management** is gathering, analyzing, and reporting on key network variables including device and link availability, throughput and utilization, and user response time.
- **Configuration management** is managing and updating network configurations including operating system versions, and network parameters (port, switch, and router settings).
- **Accounting management** is managing user and device accounts on the network.
- **Fault management** is detecting, logging, and notifying administrators of issues or faults within the network.
- **Security management** is controlling access to the network and monitoring the network traffic for security threats and breaches.

Production control engineers should make sure tools are available to help personnel to perform these functions and they are trained in their use. This usually involves working with IT organizations that already have tools and expertise. Cisco offers tools for network management including simple applications such as the Cisco Network Assistant as well as advanced applications such as CiscoWorks. One key aspect of network management that does need particular attention is IP address allotment and allocation.

IP Addresses

Establishing policy and managing the IP addresses are relevant to a control engineer. Typically, any Industrial Ethernet device (new or replacement) needs an IP address assigned to it. Many production facilities use statically assigned addresses where someone has to decide the address and configure end devices with their IP addresses. As most automation and control applications use the IP address directly in their programs, this is a straightforward way to make sure they stay in-synch, although as a facility grows, it can become a maintenance burden. Therefore, some facilities use dynamically administered IP addresses, where every time the device starts, it gets its IP address from a network service, for example using the Dynamic Host Configuration Protocol (DHCP) Option 82. The network service can be configured to issue consistent IP addresses so automation and control programs do not have to be changed, provided the appropriate network design and configuration. Lastly, control engineers should also ensure they get enough IP addresses allocated and have an allocation method that allows factory floor devices to be easily recognized. IT is usually responsible for allocating enterprise IP addresses.

Using DHCP Option 82

Ethernet switches provide excellent connectivity and performance; however, each switch is another device that must be managed on the factory floor. To make switched Ethernet networks easy to support and maintain, intelligent switches include built-in management capabilities. These intelligent features make it easy to connect manufacturing devices to the network, without creating additional configuration tasks. And they help minimize network downtime if part of the network should fail. One of the most useful intelligent features in a switched Ethernet network is Option 82.

In an Ethernet network, DHCP lets devices dynamically acquire their IP addresses from a central server. The DHCP server can be configured to give out the same address each time or generate a dynamic one from a pool of available addresses.

Because the interaction of the factory-floor devices requires specific addresses, Industrial Ethernet networks usually don't use dynamic address pools. However, static addresses can have drawbacks. Because they are linked to the MAC address of the client, and because the MAC address is often hard-coded in the network interface of the client device, the association is lost when a client device fails and needs to be replaced.

Extended fields in the DHCP packet can be filled in by the switch, indicating the location of the device requesting an IP address. The 82nd optional field, called Option 82, carries the specific port number and the MAC address of the switch that received the DHCP request. This modified request is sent on to the DHCP server. If an access server is Option 82-aware, it can use this information to formulate an IP address based on the Option 82 information. Effective use of Option 82 enables manufacturers to minimize administrative demands and maintain maximum network uptime even in the event of the failure of individual devices.

DHCP Port-Based Allocation

Another option for DHCP management for industrial automation and control system networks is DHCP port-based allocation. DHCP port-based allocation is a Cisco switch feature that simplifies IP address management and configuration in a manufacturing environment by having the switch consistently assign the same IP address to an end device.

When using the DHCP port-based allocation feature, the switch will provide the same IP address to a replaced end device when connected to the same port as the failed device. Therefore, if an end device in an industrial network

always needs to have the same IP address, this feature will allow that device to be replaced in the event of failure, upgrade, or other event, and automatically get the same IP address from the switch, without having to use network management software to configure the switch, or have applications that handle dynamic IP addresses for end devices. This eliminates the need to configure replaced devices with IP addresses by factory floor personnel.

DHCP port-based allocation is a Cisco only feature, and requires evaluation of the capabilities of end devices to support DHCP, and the overall DHCP architecture for the industrial network for effective implementation.

Conclusion

The migration to Ethernet in manufacturing environments has been growing steadily as companies recognize the many benefits that Industrial Ethernet can deliver. According to an ARC Advisory Group study, the market for Industrial Ethernet devices grew at more than a 50 percent annual rate from 2001 to 2003.

The reasons behind the success of Industrial Ethernet are clear. The technology lets manufacturers standardize and consolidate their different manufacturing network architectures, using products offered by a variety of equipment vendors. Because Industrial Ethernet is a standard technology, it enables companies to take advantage of economies of scale, while still providing the flexibility needed to support their specific factory-floor requirements. Because Industrial Ethernet uses the intelligent networking features found in corporate data Ethernet environments, organizations can enjoy substantially greater control over their networked manufacturing equipment.

A well-implemented Industrial Ethernet network can do much more than simply emulate the functions of a traditional manufacturing network. It enables companies to more closely link their internal data networks with the factory floor to make the entire company's operations more efficient. And by enabling manufacturers to tap the innovation underway that supports the millions of existing Ethernet networks, it can make possible a wide range of new applications to support business needs well into the future.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1002R)