



White Paper:

How to Develop High Reliability Ethernet Control Systems Using Media Redundancy

*Extending Ethernet Redundancy Schemes
All the Way Down to the I/O*

How to Develop High Reliability Ethernet Control Systems Using Media Redundancy

Ethernet on the Plant Floor

Ethernet was never designed for the plant floor industrial environment. It was designed for office systems and the environment of the enterprise business systems. In the business systems environment, making sure that the network has ultra high availability at all times is not necessary. For example, servers are often rebooted during off-peak use times, like at night, and on weekends. This of course will not work for the plant floor—there are few, if any, off-peak use times. A plant floor network must have availability 24/7 and probably 24/7/365. Networks cannot go down, and servers cannot be rebooted in the middle of production.

But because Ethernet is entirely ubiquitous throughout the rest of the enterprise, its cost-of-implementation and cost-of-use has been reduced to the point where it is also commonly found on the plant floor. Most IT specialists understand and can troubleshoot an Ethernet network without special training, too. So it is entirely understandable why Ethernet has become the network of preference on the plant floor as well.

The key to using Ethernet on the plant floor is to make the implementation highly available. That means in practice, various means of redundancy are employed. A redundant system must be fault tolerant. We typically make systems redundant by duplicating system components that have high probabilities of failure, like power supplies, software file systems, and other computer hardware. We also add what is called “media redundancy” which refers to the formation of a backup communication path when part of a network is suddenly unavailable. A redundant media system can survive the break of a network cable or link by switching to alternate communication paths as soon as the break is detected.

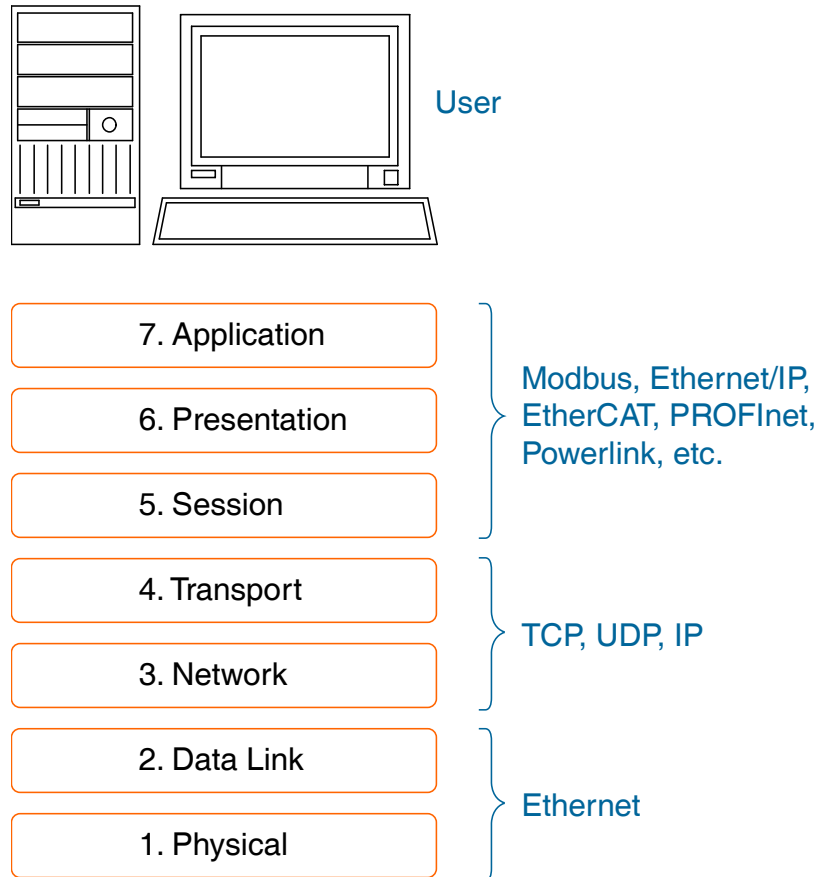
Building media redundancy into an Ethernet network isn’t simply a matter of adding another cable connection. The Ethernet devices on the network must have the ability to manage communication over duplicate paths. Ethernet normally does not allow duplicate message paths in the same network, because doing so would cause the messages to travel the network infinitely, or “loop.” This drives so-called “message storms” that can ultimately shut down the network, or prevent real communication between network devices.

Hubs, Switches and Ring Topologies

Ethernet networks began by connecting devices directly to other devices. It quickly became apparent that some means was needed to add more than two devices to a network. This device is called a hub, because the devices were connected in a set of spokes to the hub. This sort of configuration is also called a “star” topology. More than one “star” can be connected through another hub to form a “tree.” Star topologies are very practical for connecting edge devices, such as sensors, scanners, and other I/O devices to the network, but they do not address redundant message paths well.

A network hub electrically repeats an incoming network message at all of its connection ports. By itself, a network hub cannot be made redundant, because it has no way of knowing which messages are redundant, and therefore passes all of them. This can cause packets to collide and become garbled as messages are received and reformed. “Packet” or “message” storms can be created that could effectively shut down the network.

Figure 1: OSI Reference Model



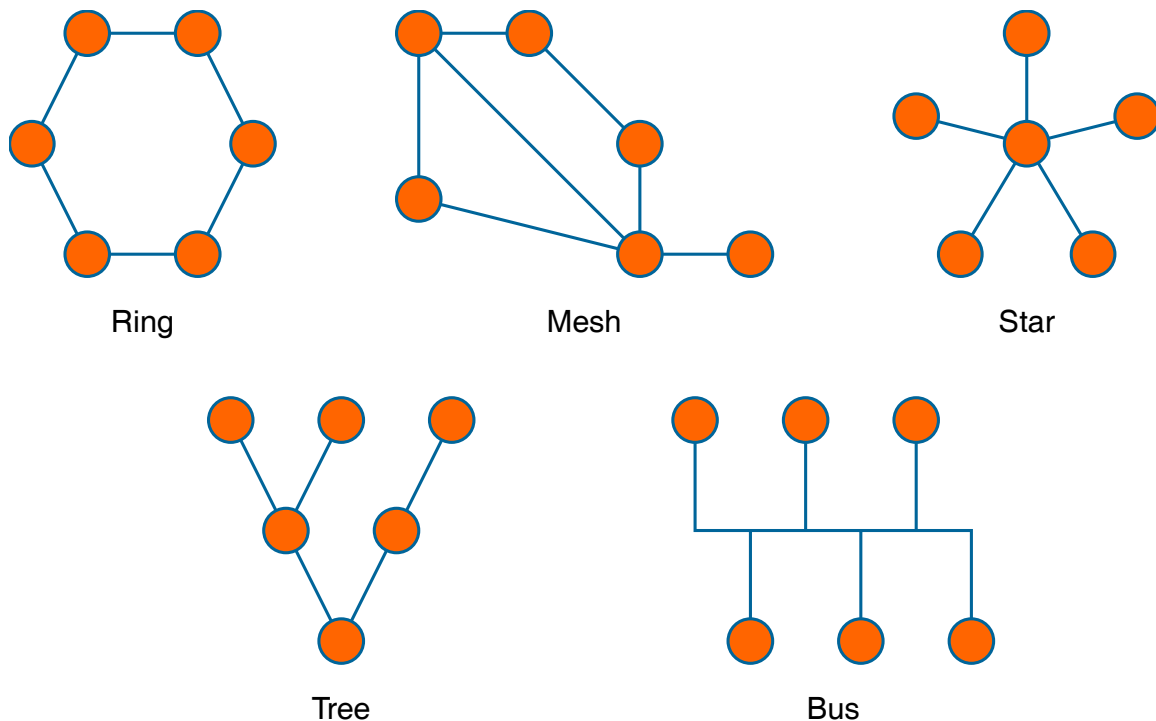
The Seven-Layer OSI Networking Model.
Ethernet defines Layers 1 and 2 of the network topology.

Hubs work at the lower layer of the OSI Reference Model, the physical layer. Other devices, called bridges by the IEEE, and “switches” by most everybody else, perform the same function as a hub, but with more intelligence, because they sit on the second or data link layer of the OSI model. Switches are similar to hubs, however, they learn which of their ports are connected to which devices, and then forward messages to the correct port only, effectively reducing network traffic and preserving network bandwidth. Switches are designed to perform other intelligent functions to filter and direct network traffic as required.

Switches can be “cascaded” or added upstream or downstream of switches already in the network to further segment the Ethernet network.

Before we discuss how switches can be used to manage redundant message paths in a network, we should review the five basic topologies used to form Ethernet networks: Star, Tree, Bus, Ring and Mesh. Bus networks have a single cable backbone, with many devices connected to it. Foundation Fieldbus HSE is an example of an Ethernet network formed as a Bus. The single cable backbone is the common point of failure, so industrial plant bus networks are usually designed with dual redundant cable pairs. Star topologies have all the nodes connected to a central site, or hub. It has the advantage over the Bus that if cable breaks between the hub and any node, the operation of the other nodes on the network are not impaired. Tree topology is a combination of Star and Bus topologies, providing the ability to link together smaller networks or Stars via a common Bus. Mesh topologies use lots of wire to connect every node with every other node, providing alternate paths for every node to maintain communication in the event of failure. Mesh topologies are very expensive, but are very fault tolerant. It is worth noting that the Internet is itself a mesh network, designed that way to operate even in the event of damage during a nuclear war. The last of the five is the Ring. In this topology, devices are connected in a series ring and messages go only one way from device to device.

Figure 2 : Basic network topologies



With the exception of the more expensive Mesh topology, none of these ways of designing an Ethernet network provides any means for handling redundancy or redundant message paths. Recognize that Bus and Ring topologies can have single point of failure problems.

Ethernet Redundancy Schemes

So how do we build redundant Ethernet networks without causing message loops? We simply utilize switches that can manage multiple connection paths to any one device.


That is, we connect our devices with switches capable of detecting a redundant path to the same address. This works independent of topology. There are four of these methodologies, or protocols, for achieving redundancy with switches: Spanning Tree (STP), Rapid Spanning Tree (RSTP), Proprietary Ring and Trunking. Each of these redundancy methods will take care of media failure within the network structure itself, that is, between the switches and hubs, but not necessarily to the network devices they connect to.

Spanning Tree Protocol, or STP was the first protocol to be developed based on the IEEE 802.1D Media Access Control (MAC). Spanning Tree is based on timers and will operate with link segments and coaxial cable segments. Rapid Spanning Tree is a modernization of the STP protocol which, as its name implies, operates significantly faster than its older predecessor. Both protocols aren't limited to Ring topologies, but will also work well with a Mesh topology.

The Proprietary Ring is a combination of a topology and proprietary protocol that permits a network Ring to continue to function when one of its segments is broken. A master switch is set up to monitor and control packet traffic in a proprietary way. Proprietary rings are simple to understand and set up, but they have the disadvantage of being proprietary, and this limits hardware selection to a single vendor for the entire network. The Proprietary Ring is very fast, recovers well, and does an excellent job of managing redundancy with minimal usage of switches and cabling.

Trunking, which is sometimes called link aggregation (as it is in the IEEE 802.1 standard), provides two or more parallel paths between device ports for redundancy. Trunking has the advantage of increased bandwidth and throughput, since it provides dual paths for data transmission between each switch.

Figure 3: Trunking Diagram



Fully Redundant Ethernet Networks Means All the Way to the I/O

Because of Ethernet's heritage from the enterprise or business system networks, most network designers limit themselves to containing media failure within the network infrastructure itself. That means that they provide redundant connectivity between the hubs and switches of the network, but they generally omit the devices the network connects to. In the case of plant floor networks, that generally means the I/O.

Ethernet I/O comes in two basic flavors. Some devices have digital connectivity themselves, such as the network ports found on barcode scanners or weigh scales. Other devices act as I/O aggregators with multiple analog, discrete, or digital inputs and outputs.

Figure 4: Types of I/O



Ethernet I/O aggregators are available in a variety of form factors. Acromag's BusWorks[®] and EtherStax[®] series demonstrate configurations to interface just a few or very many analog and discrete I/O signals.

In order to provide appropriate redundancy to the network, we have to consider what the purpose of redundancy in a plant floor network is: high availability of the control system. That means we must provide redundancy not only to the network topology, but also all the way to the I/O, or we continue to have single point of failure issues that may cause the control system itself to fail.

Back to the Hub

Remember that a hub will repeat an incoming message at all of its other ports. If two ports of the hub happen to be connected to an external redundant network switch, this causes the switch to immediately detect the redundant path, disable it, and hold it as a back-up path should the primary path fail. If the primary path fails, it will trigger the fail-over and recovery operation using the particular protocol of the redundant network switch. This provides some significant advantages when multi-port hub functionality is combined and embedded inside an I/O device. This simple methodology as implemented by Acromag's EtherStax I/O system allows the I/O to become interoperable with standard Ethernet redundancy protocols like STP or RSTP. Further, it allows the I/O to also work seamlessly inside any proprietary Ring network. This means that media redundancy and reliability are now easily implemented all the way back to the "end-node" device rather than stopping at the connected switch.


The Ethernet ports on the I/O can operate as switches or hubs which allows their use in virtually any Ethernet network architecture or in any redundancy scheme supported by switch manufacturers. Designers and users win. It is a simple, cost effective, and reliable approach which can also provide a future upgrade path based on standards, so that lifecycle costs are minimized.

Methodology for Fully Redundant Ethernet Networks in Control Systems

The premise for fully redundant Ethernet networks in control systems is to use devices with dual Ethernet ports that can emulate the behavior of a hub. If an I/O or control device is able to operate its dual network ports as a hub, this easily brings media redundancy down to the end-node.


In its simplest form, a Redundant Media Path is formed with an Ethernet I/O device, such as an EtherStax unit, that is connected to two ports of the same redundant switch. This switch would use a Spanning Tree Protocol to manage redundancy.

Figure 5: Simple Node Redundancy Connections (One RSTP Switch)



A break in either path to the unit will recover communication on the opposite path in about 15 seconds or less. With RSTP switches, recovery times can be much faster, perhaps 1 or 2 seconds. Additional redundant paths can be added by simply installing a second redundant switch.

Figure 6: Simple Node Redundancy Connections (Two Switches)



Because of the hub functionality, redundancy to the I/O functions independently of the system topology and can accommodate mesh, ring and tree configurations.

Figure 7: RSTP Mesh Connections and Ring Redundancy Connections






Figure 8: Redundant Ring Network with a Wireless Link



The second port of the Ethernet I/O devices may also function as a backup path for other types of connection media, such as Ethernet-to-Wireless radio transmission, as shown above. You could just as easily replace the radio path with an Ethernet serial server, Ethernet-to-Cellular or Ethernet-to-Satellite Modem to form a backup path along some other media as well. In the case shown, the primary path is the hard-wired copper connections for “mission critical” operation and the wireless path is held in reserve as a backup.

Conclusion

Which method you ultimately choose to manage redundant media should consider future upgradeability and expansion, the recovery time, its flexibility in wiring, and the amount of switch ports and cable that will be consumed.

| | STP | RSTP | RING | TRUNKING |
|--|---------------|---------------|-------------|------------------|
| Protocol Reference | IEEE 802.1D | IEEE 802.1W | Proprietary | Proprietary |
| Recovery Time | 30-60s | 1-2s | ≤ 300ms | ≤10ms |
| Wiring Flexibility | Very Flexible | Very Flexible | Restrictive | Very Restrictive |
| Switch Supplier Flexibility | High | High | Low | Low |
| Amount of Cable Needed for Media Redundancy | High | High | Low | Highest |
| # Ports Consumed | High | High | Low | Highest |
| # of Switches Needed | High | High | Low | Highest |

With RSTP, the Recovery Time for a media failure between switches ranges from 1-2 seconds, while the recovery time to the Ethernet I/O may be longer depending on how the switch software parameters are configured (e.g. Max Age Time, Hello Time, Forward Delay Time, etc.). For redundancy applications where general recovery times of several seconds or longer are acceptable, then using switches with standard RSTP protocol offers the most flexibility and interoperability among switch vendors, although at a cost of using additional switches and cable.

For control or I/O system applications where recovery times are more critical and costs may be sensitive, Ring Redundancy offers media recovery times of less than 300 milliseconds for any ring path. Further, when using Ring Redundancy combined with Ethernet I/O such as EtherStax, this methodology offers “end-node” redundancy at a faster rate and lowest cost based on the number of switches and cables needed.

Trunking, or Link Aggregation as mentioned earlier, is a method that uses multiple paths between switches to achieve the greatest bandwidth and more redundancy levels. With recovery times of less than 10mS, it is currently the fastest media redundancy method. It is ideal for high-speed redundancy between servers and information systems. Since trunking is proprietary and costs more due to its restrictive nature, trunking is not considered a viable or practical redundancy protocol for end-node industrial i/o or control system applications.

In conclusion, by integrating hub functionality into end devices, it is therefore possible to extend standard and open Ethernet redundancy methodologies all the way to the I/O. The end results allow users to utilize “off the shelf” Ethernet switch technology to provide the very high availability and reliability necessary to operate in an industrial plant floor environment.

NOTE: Recovery times listed are estimates based on bench tests. Actual times may vary.

About the Author:

Acromag is an international corporation that has been manufacturing and developing measurement and control products for more than 50 years. Acromag offers a complete line of industrial I/O products including process instruments, signal conditioners, distributed I/O systems, network communication devices, and data acquisition boards.

For more information about Acromag products, call the Inside Sales Department at (248) 295-0880, Fax (248) 624-9234, or e-mail sales@acromag.com. Write Acromag at P.O. Box 437, Wixom, MI 48393-7037 USA. Our web site address is <http://www.acromag.com>.

BusWorks and EtherStax are registered trademarks of Acromag, Inc.