

AREAS FOR REDUNDANCY IN ETHERNET/IP™ SYSTEMS WITH FOCUS ON ETHERNET NETWORKS

Dominic Iadonisi
Industrial Business Development Manager
Ruggedcom Inc.

Presented at the ODVA
2007 CIP Networks Conference & 12th Annual Meeting
September 18-20, 2007
Englewood, Colorado USA

Abstract

The need for redundancy for EtherNet/IP is centered on the concept of maximum up time. If you have a redundant system, it can take hits to operations and still maintain full functionality. Redundancy should not come at the price of extreme costliness and a poor return on investment. Redundancy in EtherNet/IP networks cover various levels: Physical, Data Link and Network (IP). Physical refers to the actual hardware and cabling, Data Link refers to the data interconnections at Layer 2 of the OSI model. Network refers to Layer 3 (IP) of the OSI model and the Routing Communication protocols that provide redundancy at the IP layer. Redundancy within EtherNet/IP also comes at a price. The more redundancy built into a system the higher the cost of the system. So it becomes a balance of amount of redundancy, maintainability and cost. Helping the customer determine what is needed and providing the tools to make educated choices is key. This paper explores the different aspects of redundancy in respect to EtherNet/IP based systems and helps the end customer determine the level of redundancy necessary to maintain their required level of operation for profitability.

Keywords

EtherNet/IP, CIP, Ethernet, Physical, Data Link, Network, Redundancy, Cost, Balance

Introduction

With EtherNet/IP, as with any Ethernet based Industrial protocol, redundancy is needed to maintain maximum uptime and still be able to deal with minor outages and failures to the environment. This all gets rolled into the reliability of the entire system, from the very edge devices, through the network core, to the plant backbone.

The definition of **Reliability**, according to Webster's is:

noun

- the quality of being dependable or reliable [syn: dependability]
- the extent to which an experiment, test, or measuring procedure yields the same results on repeated trials

The definition of **Redundancy** (in relation to electronics and systems), according to Webster's is:

- Duplication or repetition of elements in electronic equipment to provide alternative functional channels in case of failure.
- Repetition of parts or all of a message to circumvent transmission errors.

The entire system, from edge devices to the network and communications should be held to the same set of standards for optimum uptime and supportability.

Industrial System Reliability covers multiple areas-

- Robustness of the Application
- Handling Environmental Stresses
- Data handling and path redundancy
- Data integrity

Table 1 identifies the main differences between Commercial or Enterprise type networks and Industrial rated networks. It is important to keep these in mind when looking into building an EtherNet/IP network.

	Enterprise	Industrial
Installation	<ul style="list-style-type: none"> ▪ Fixed basic installation in the building ▪ Typically use RJ45 connectors kept "in the clear" ▪ Network is normally centrally located in wiring closets 	<ul style="list-style-type: none"> ▪ Plant dependent cabling and cable ducting ▪ Field attachable connectors with ratings up to IP67 ▪ Network is distributed throughout the plant in panels and small cabinets
Data	<ul style="list-style-type: none"> ▪ Medium to Large data packets ▪ Medium to high network availability ▪ Predominantly acyclic transmission-file downloads, printing, Internet, etc. ▪ Real-time behavior not necessary 	<ul style="list-style-type: none"> ▪ Small data packets ▪ Very high network availability ▪ Predominantly cyclic transmission-know data rate with programmed request/reply intervals ▪ Real-time behaviour necessary, especially where coordinated motion is used (IEEE 1588)
Environment	<ul style="list-style-type: none"> ▪ Normal temperature range ▪ Little dust, moisture and vibration ▪ Hardly any mechanical loads or problems with chemicals ▪ Low EMI/RFI/EMC Requirements 	<ul style="list-style-type: none"> ▪ Extended temperature ranges- very low to very high ▪ Dust, moisture and vibration likely ▪ Risk of mechanical damage or problems with chemicals in hazardous environments ▪ High EMI/RFI/EMC requirements

Table 1: Differences between Enterprise and Industrial Networks

Table 1 lists some of the requirements and conditions industrial networks require equipment that is optimized for the environment that it is installed in. This can vary widely from Deserts in the Middle East to the Arctic Circle.

The three main areas of coverage for Ethernet-based industrial systems redundancy are physical, data link and network as show in Figure 1 below. We will be tackling each area individually. We also will be looking at the balance between the cost of making a system redundant and the cost of failure of the system and lost production. Security is a separate topic and will not be handled in this document.

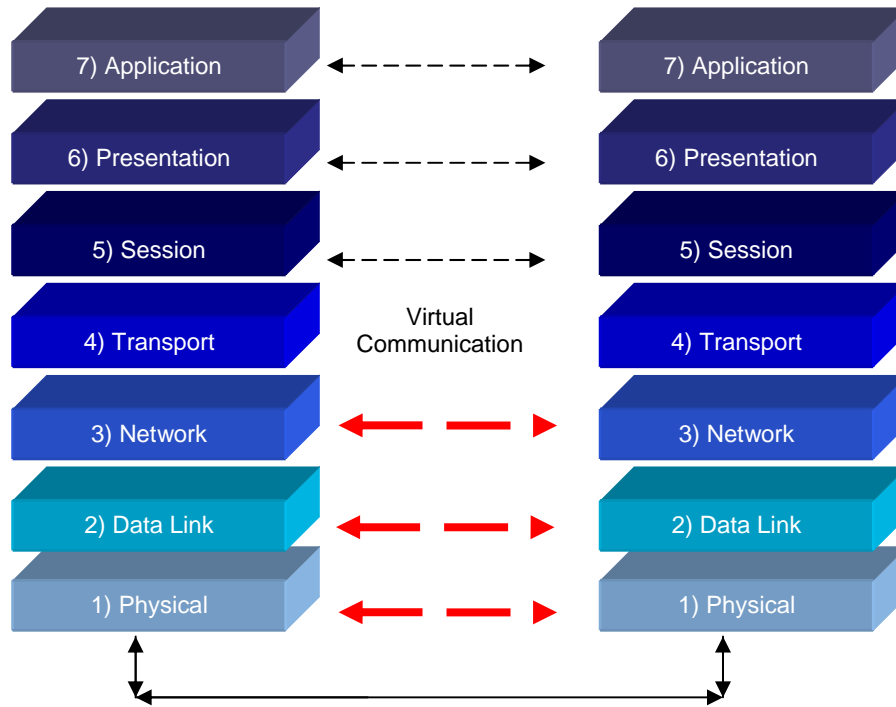
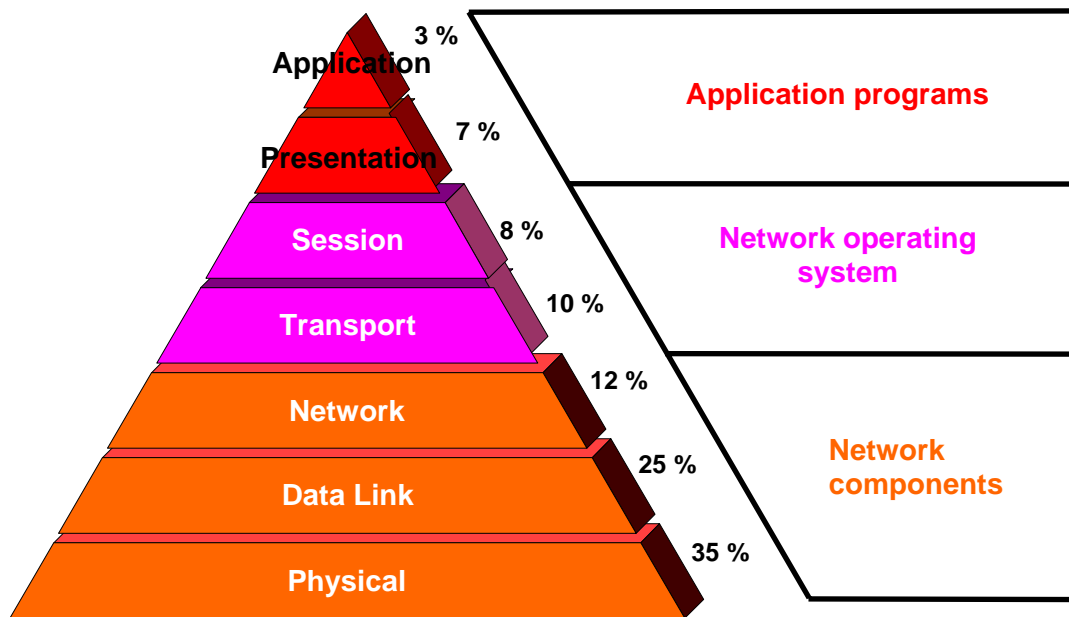


Figure 1: Areas of redundancy focus in the OSI model

The lower you look at the OSI model, the greater the impact of failures that occur. For example, if you lose a cable connecting an end device to a switch port, there is no data movement of any kind and there is an impact for the entire process, depending on the importance of that end device. If there is an issue at layer 3 where a router may have experienced a loss of power and loses connection to the plant backbone, localized plant processes can still operate, but operation for distributed plant or business processes can be effected to a large extent.

Figure 2 shows impact upon the network based upon percentages. Again, notice that the lower you go, the more impact failures have, with 72% of failures occurring on the first 3 layers. These include hardware failures, cabling failures, power losses, programming miss-configurations, etc.



Source: Datacom, Network Management Special

Figure 2: Area of System failure with percentages

Section 1

Physical redundancy: More than just the cabling

Physical redundancy covers the physical Ethernet network connections (and Ethernet equipment) AND the physical hardware the connections go between. Network redundancy focuses upon the multiple routes that can be used between edge devices. The more available routes edge to edge, the more failures the network can sustain and still keep the process alive and functioning. Physical redundancy normally follows two scenarios:

1. Diverse routing of cabling – if a cable tray or conduit is damaged in some way cutting link on the cabling in it, there is another way to get where you need to by providing cabling by another route to maintain connection reliability.
2. Redundant hardware – having multiple connections on a controller or other hardware allows the controller reliable connectivity in cases where a connection or port has suffered a failure. This also involves having redundant network hardware in case of failures, including multiple power supplies, multiple CPU cards on controllers, etc.

When looking at redundancy for the EtherNet/IP network, you need to first look at the application and the area of coverage, including the number of devices that are attaching to the Ethernet network. Are they grouped according to location and function? Application performed? Device type? Will there be a requirement to connect to the existing Plant Backbone network? Based on this knowledge you can begin to put together a picture of how the Ethernet network will look and what the number of ports will be on the Ethernet switches that will be put in those areas. This is needed to determine number of cables, physical routing of the cables, and location of network nodes to connect the cables, and so on.

A popular way to look at system connectivity needs is the Zone/Cell view where you have a Zone of control divided up into functional Cells. Refer to Figure 2 above.

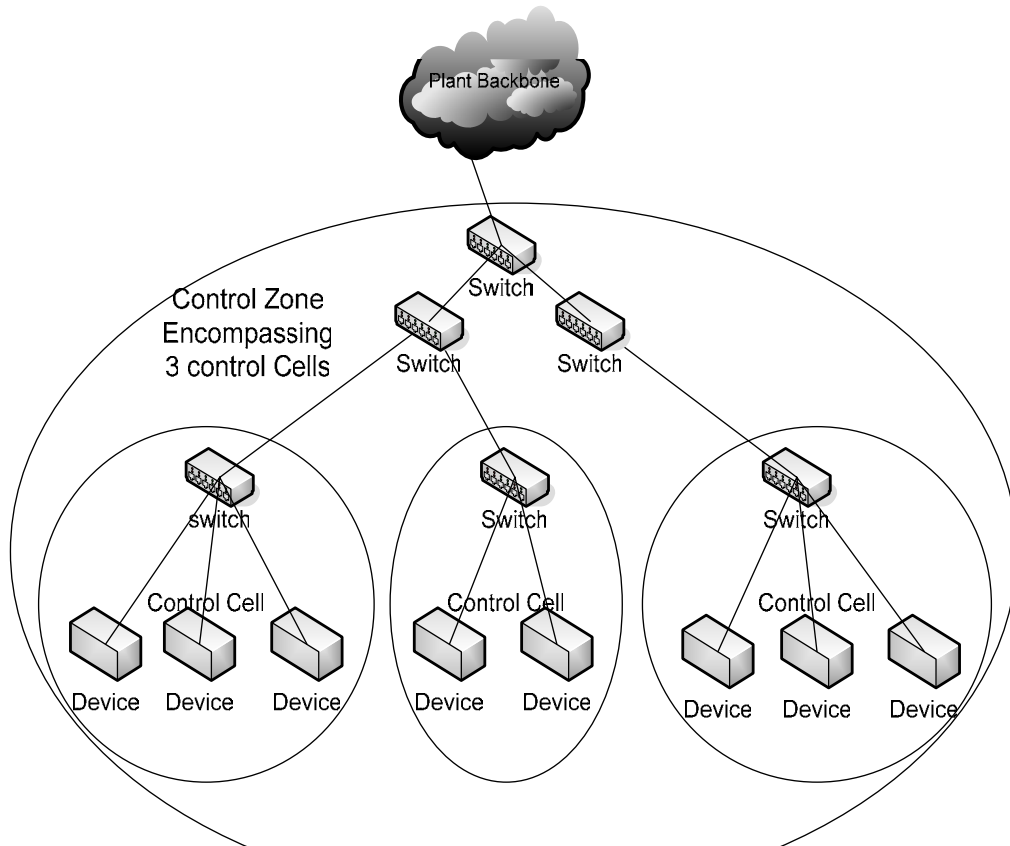


Figure 3: Control Zone/Cell Reference Diagram

Looking at Figure 3, assuming that each line is a single cable connection, it would be very easy to isolate sections of the process with the loss of just 1 or 2 cables. At the physical layer, it is important to plan out redundant connections to devices that can support multiple connections. Many devices only have one data interface, but the Ethernet switches they connect to have multiple ports to support connections to other switches, forming redundant paths and being able to work around port and cabling failures. In the next sections we will discuss what network protocols are available to make the best use of these redundant paths between network nodes.

Once you have decided how the devices are connected to the network, you then have to decide the level of redundancy needed for the maximum expected uptime of the system. This requires evaluating the cabling and Ethernet network hardware needs. Do we need redundant cabling between devices? Is running redundant cable by different routes needed to provide physical security of the cabling in case of damage to one of the runs? Is there more than one Ethernet interface on the device to be used (many controllers have multiple Ethernet interfaces in case of Ethernet port or module failure)?




Figure 4: Ethernet Network Physical Topologies

Section 2

Data Link Layer: Using the Ethernet switches in the network to provide protocol redundancy and maintaining Ethernet network health

Layer 2 Redundancy protocols do two things: Identify all the possible paths amongst the networking devices and place the redundant extra paths in a blocking state to remove network loops. Loops in an Ethernet network cause data duplication and will bring a network to its knees in a short period of time. In the event that the network segment fails, the protocol activates the appropriate ports that are in a blocking state to reestablish connectivity. The object being to fix the issue before the process even knows there is a problem.

Ethernet networks have redundancy protocols that are supported by identified Ethernet standards. These are supported in Layer 2 and Layer 3 of the OSI model. First we will look at Layer 2:

Standard Layer 2 Network Redundancy Protocols

1. Spanning Tree –There are several flavors of Spanning Tree:
 - a. STP (Spanning Tree Protocol) – Standardized in 1996 as IEEE 802.1D, it is the first and slowest of the Spanning Tree protocols. Average failover time for STP started at 30 seconds and went up. Way too slow for any industrial Process.
 - b. RSTP (Rapid Spanning Tree Protocol) – Standardized in 1998 as IEEE 802.1w, it was an evolutionary leap for STP. It is more rapid, with failover times from about 500msec to up to 12 seconds, so it was better than STP. There still remains an issue with the speed of failover for Industrial processes.
 - c. MSTP (Multiple Spanning Tree Protocol) – Originally standardized as IEEE 802.1s and then incorporated into IEEE 802.1Q 2003, it allows multiple instances of Spanning Tree Protocol per Virtual LAN. This means that in a single physical network, there can be multiple virtual network groupings, each with their own instance of Spanning Tree Protocol.

- d. There are proprietary implementations of Spanning Tree that are optimized for use in Industrial Networks. They are based upon standard RSTP, but are not designated as a standard STP protocol.
2. LACP (Link Aggregation Control Protocol) – This protocol allows the user to configure multiple Ethernet ports between Ethernet switches into a Single virtual “Link”. This allows load sharing of information between the links and is extremely fast in moving data between a failed port and an adjacent port if there is a link failure.
 3. The amount in interconnections amongst the network elements dictates the amount of failures the network can take and still maintain the process.

Figures 5 and 8 show examples of the protocols.

Spanning Tree is a redundant topology in that it provides network redundancy instead of just path redundancy while preventing loops in a network. For Ethernet to function properly only one active path can exist between devices. To provide redundancy, Spanning Tree relies on having multiple paths or connections to different switches and configures some of these paths into a standby (Blocked) state. If a network segment becomes unreachable, spanning tree reconfigures and reestablishes link by activating the "Blocked" links.

All switches in the LAN gather information about each other through an exchange of data messages called BPDU's or Bridge Protocol Data Units. The exchange of messages causes the following:

- The election of a "Root" switch for stability.
- The election of a designated switch.
- The removal of loops by placing redundant switch ports in a backup state.

The "Root" switch is considered to be the "logical" center of the Spanning tree network. All paths that are not needed to reach the "Root" switch from anywhere in the network are placed in backup mode. BPDU's contain information about the transmitting switch it came from and it's ports including:

- Unique switch Identifier or MAC address.
- Switch priority
- Port priority
- Port cost.

Spanning Tree uses this information to elect the "Root" switch and "Root" port for the switched network. The switches send configuration BPDU's to configure the spanning tree topology. All switches connected to the LAN receive the transmitted BPDU. The BPDU's are not forwarded by the switch, but the information contained in the BPDU can be used by the receiving switch to transmit a new BPDU.

The resulting action of this communication is:

- One switch is identified as the Root.
- The shortest distance to the root is determined for each switch.
- A designated switch or switch closest to the Root is selected.
- An active port from each switch is selected and the others are blocking.

If all the switches are enabled with default settings, the switch with the lowest MAC address becomes the root by default. However, due to traffic patterns, number of forwarding ports or just simply physical location, this may not be the best switch to be the root. By increasing the priority (lowering the actual numerical value of the priority number) of the ideal switch so that it becomes the root, you are forcing spanning tree to recalculate and form a new topology. The same can be said for which port is active and which port stays in standby. By increasing the priority (lowering the actual numerical value of the priority

number) of the ideal port so that it becomes active, you are forcing spanning tree to recalculate and form a new topology.

Each port on a switch using spanning tree protocol in one of five states:

- Blocking
- Listening
- Learning
- Forwarding
- Disabled

Each port moves through these five states as follows:

- From Initialization to blocking
- From Blocking to Listening or disabled
- From Listening to Learning or to Disabled
- From Learning to Forwarding or to Disabled
- From Forwarding to Disabled

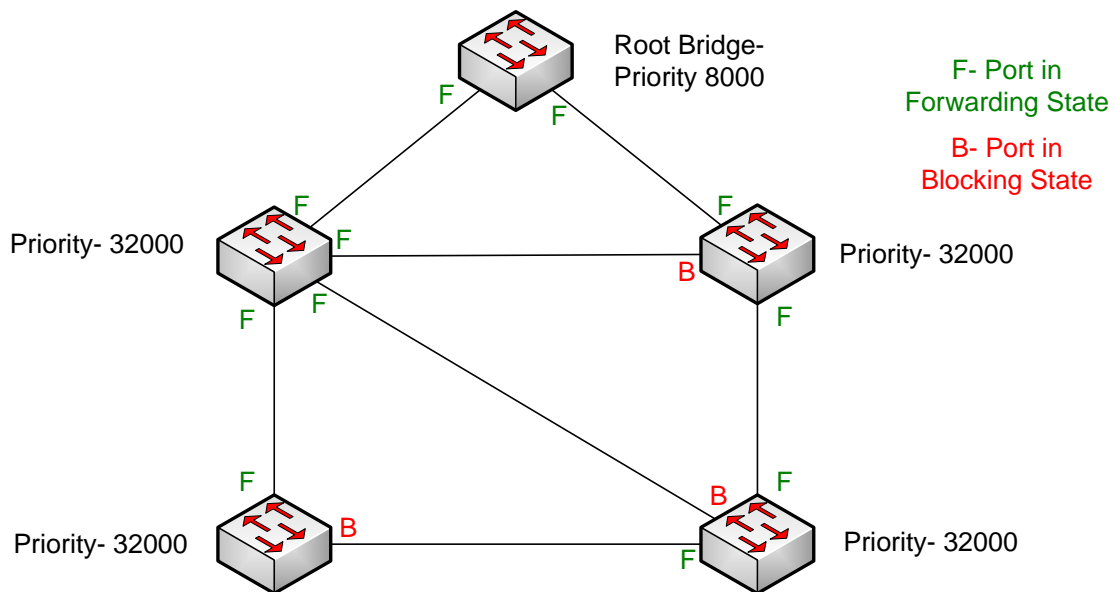


Figure 5: Example of a Spanning Tree Ethernet Network

Spanning Tree networks can support either ring or mesh topologies. A ring topology is basically a ring of Ethernet Switches connected together in a ring fashion. A mesh topology requires the use of a couple of Ethernet switches up at the top with switches below that have connection to both the upper switches. Mesh networks use more fiber than ring networks, but can typically survive more network hits intact. Figure 5 shows a typical Mesh network while Figure 6 shows a Ring network example.

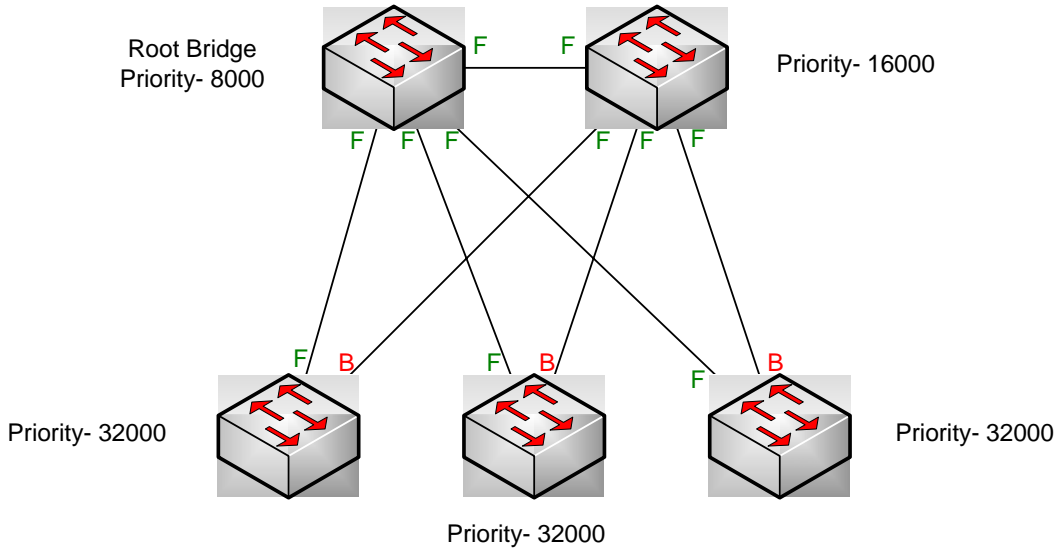


Figure 6: Spanning Tree in a Mesh Network

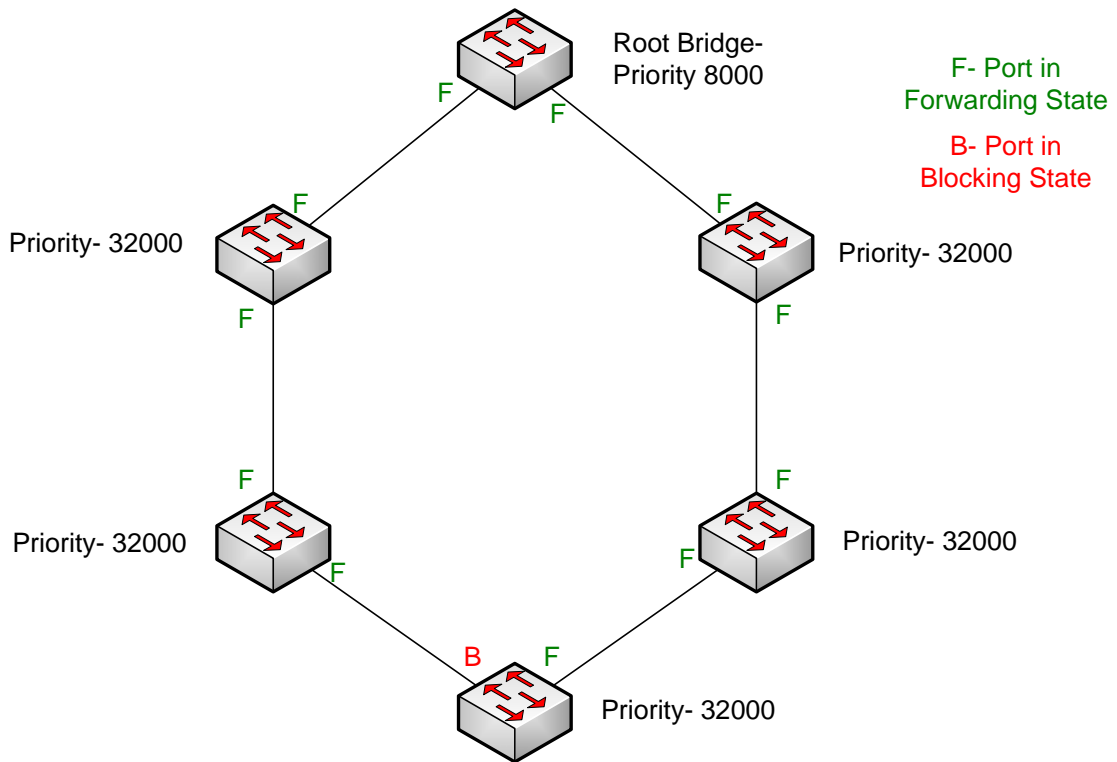


Figure7: Spanning Tree in a Ring

Link Aggregation Control Protocol (IEEE 802.1ad) provides redundancy without the use of Spanning Tree. It enables users to be able to bundle groups of ports between switches to form 1 virtual link with the bandwidth of the member links. LACP provides several functions:

- A. Higher bandwidth
- B. Enhanced Bandwidth Granularity
- C. Load sharing across the member links to balance bandwidth across the member links
- D. Fault tolerance provided by offloading data to working member links when a member link fails

LACP is a method of providing needed extra bandwidth between Ethernet switches that have extra non-utilized ports without buying a switch or switches with higher bandwidth ports. For example, moving from 100Mbps switching to Gigabit Ethernet switches.




Figure 8: Example of an LACP based Ethernet connection between switches

Non-Standard Layer 2 Redundancy Protocols

There are many identified network redundancy protocols that are not standard and vendor specific that are designed to provide a very fast fail-over mechanism for the controls network. These protocols typically provide faster recovery than Spanning Tree protocols. But, they are not standard. This can make interoperability difficult amongst different systems and networks that may contain differing vendor's products. They use ring architectures and the breaking of the ring to prevent loops is through the use of a "Redundancy Manager". The Redundancy manager breaks the ring by placing one of its ring ports into a blocking state. If a link is broken in the ring, the blocked port is placed into a forwarding state to make sure that the network connectivity is maintained. If more than one link is broken, then ring segments become isolated until the broken links are fixed. Figure 9 shows an example of this ring topology.

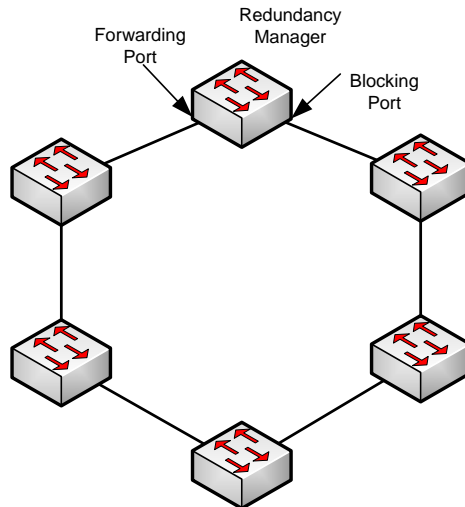


Figure 9: Example of Non-Standard Redundancy Protocol Ring Topology

Section 3

Network Layer Redundancy Protocols – How Routers talk to each other and fix breaks

It is growing more and more apparent that as EtherNet/IP networks expand, the use of a single IP subnet is not enough. In order to facilitate communication between IP Subnets, you need to use a Layer 3 network device, namely, a router. Routers can provide data movement in 2 ways: Statically via routes that are mapped by hand (**Static Routing**) or dynamically via designated routing protocols (**Dynamic Routing**). Static routing can be useful for small routing areas, but does not provide fast failover because it requires user interaction to program and alternate route manually. Dynamic routing is required where a hand off failover is required or the routing environment is large. Routing protocols are inherently slower on failover than layer 2 protocols.

Routers support several types of protocols to communication like OSPF (Open Shortest Path First) and RIP (Routing Information Protocol) that have a communications redundancy built in as long as the physical network architecture remains in place.

There is also a router redundancy protocol that supports redundant router replacement. If one router fails, its designated backup is placed into service seamlessly as if the original never left. This is called VRRP, Virtual Router Redundancy Protocol.

Distance Vector vs. Link State Routing Protocols

Distance vector

- Sends routing table info only to neighbors, so change communication may need one min/router
- Also called “routing by rumor”
- Easy to configure, but slow

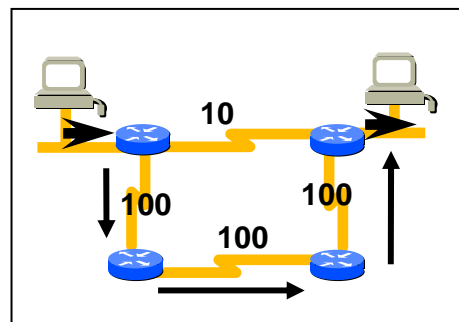
Link state

- Floods routing information about itself to all nodes, so changes are known immediately
- Efficient, but complex to configure

OSPF and RIP - Standard Router Communications Protocols

OSPF and RIP protocols are used as ways that routers communication with each other and tell each other what IP Subnets they have attached. By sending these Routing table update to each other the routers build a map of how the network is constructed at Layer 3. This also identifies the redundant ways that these routers can maintain connection to each other if a router loses connection on a port. If a router knows of another way to get to an IP Subnet it needs to send data to, it will use these alternate paths. Figures 10 and 11 show some examples of routing protocols.


OSPF is referred to as a Link-State Routing protocol. The best routes from router to router are based upon the A class of routing algorithms in which each router broadcasts connection information to all other routers on an internetwork. This saves the routers from checking for available routes but adds the memory requirement of storing all the routing information. This algorithm relies upon the cost of the links between routers, not the number of hops. The cheaper the cost on a connection indicates a higher bandwidth capability. OSPF keeps in memory ALL of the possible routes, not just the active routes.




OSPF
Industry standard that selects the path with the lowest cost

Figure 10: OSPF Routing Protocol example

RIP and RIP II are types of Distance Vector protocols. Distance vector algorithms compute distances from a node by finding paths to all adjacent nodes and by using the information these nodes have about continuing on the paths adjacent to them, router hop by router hop. Distance vector algorithms can be computationally intensive, a problem that is alleviated somewhat by defining different routing levels. They rely upon the number of hops in a particular direction between the source router and destination router. They do not take into consideration the speed of the physical media, so it is possible to move traffic across a suboptimal link



RIP
Industry standard that selects the path with the fewest hops



RIP II
RIP II has the ability to select the faster path (using load, distance, etc.)

Figure 11: RIP and RIP 2 routing protocol examples

Aspect	OSPF	RIP
Topology	Hierarchical	Flat
Memory & CPU requirements	High	Moderate
Routing table size	Large	Moderate
Controlling body	Industry standard	Industry Standard
Convergence	Fast	Slow
Configuration	Difficult	Easy
Supported protocols	IP	IP

Table 2: OSPF and RIP Comparison

Router Redundancy

VRRP is the way for routers to perform physical redundancy to each other. If one router dies or is unable to function in the appropriate manner, its designated backup will take over the former routers function. They maintain this relationship through the use of HELLO packets and regular updates to make sure that both routers have all the same information. The use of VRRP would be a function to incorporate into an EtherNet/IP design if there is a requirement to attach to a corporate network and there is a requirement to maintain some sort of segregation between the plant floor EtherNet/IP network and the corporate environment. Figure12 shows an example of VRRP.

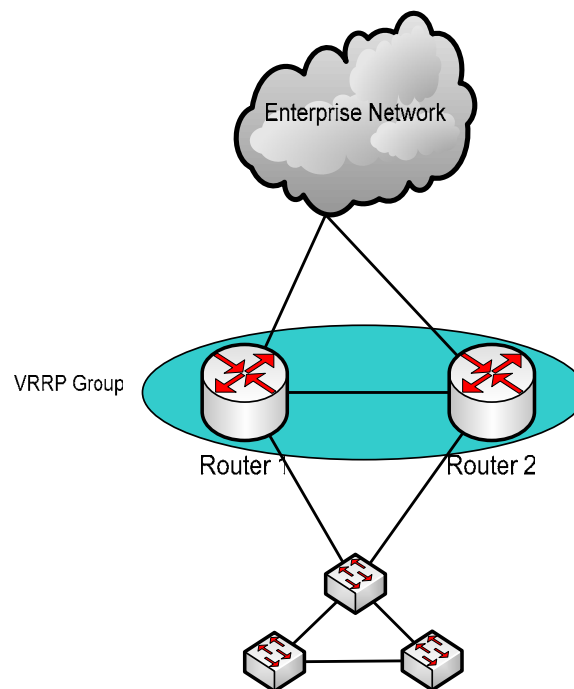


Figure 12: VRRP Example

Section 4

Determining the Cost of Redundancy: How much is too much?

Designing redundancy into a system is always a balancing act. You must think about how much to incorporate into the various areas, physical, network and application. The first thing that has to be determined is the scope of the system being installed. The following are a list of questions to ask when evaluating EtherNet/IP design:

1. Is this a new install or an upgrade to a previous installation?
2. Is there any existing cable that can be reused?
3. Is there any existing equipment that can be reused?
4. Has the area of the Installation been determined?
5. Copper or Fiber Optics? This is dependent upon distance and the environment for the installation.
6. Who is the Control System vendor?
7. Will there be a point of connection to the existing plant network? What sort of data is intended to be passed to this network from the plant floor?
8. To what extent has redundancy been considered? Is the network a ring or mesh based network?
9. If Ethernet Network redundancy is not being considered, is it economically feasible to do without it? How many outages are you prepared to pay for in lost revenue? Balance this against the cost of managed vs. unmanaged switches or more advanced Ethernet networking devices like routers.
10. How experienced are the plant controls support people in regards to Ethernet networking and will the IT staff be involved in the support?
11. What is the projected budget for the control system, including the network cabling and equipment?

Redundancy levels are dependent upon the operation expectation of the EtherNet/IP Control System being installed. Discrete automation systems usually incorporate most of the redundancy into the Ethernet network, requiring Ethernet network devices that are smarter and more expensive but able to heal around network breaks. Process control systems rely upon the controllers for redundancy, meaning the Ethernet network itself is relatively dumb, but there is double hardware expense due to use of parallel, non-redundant networks.

Based upon years of observations regarding Ethernet networks and the types of devices purchased to make them (namely buying unmanaged switches instead of Managed), it is very typical to actually see the cost differential between managed and unmanaged Ethernet switches exceeded by the lost revenue of an extended downtime event caused by a network outage. The ability to be able to monitor a network and see the application in action can help predict events that can cause outages. An unmanaged switch is in effect a “Blind” switch. It is not possible to see how the network is performing and perform predictive maintenance based upon what you can “see”. Also, the ability to use port mirroring on a managed switch can assist with troubleshooting Application Level issues as you can use a protocol analyzer to see the EtherNet/IP application in operation.

Interestingly enough, you can go overboard on redundancy as well. Using too many connections between Ethernet switches can cause slow downs in reconvergence of a network if there is a loss link or switch. Ring topologies typically use 2 interswitch links per switch, mesh topologies can use 3 or more. The recommend norm is no more than 3 for edge switches in a mesh network environment. Buying Ethernet switching hardware that exceeds the requirements for the network can cause cost increases as well.

Summary

Understanding the relationships between the physical structure of a network and the protocols that run on the network is key to creating a truly maintainable and adaptable network that deals with issues effectively. It is best to consult with the Ethernet switch vendor. That is the basis of the network being installed. Use their experience to help determine how much is needed to make your EtherNet/IP based control system a success in operation over its lifetime.

DeviceNet, DeviceNet Safety, CIP, CIP Motion, CompoNet, CIP Safety and CIP Sync are trademarks of ODVA. EtherNet/IP is a trademark of ControlNet International under license by ODVA. Other trademarks are property of their respective owners.

The ideas, opinions and recommendations expressed herein are intended to describe concepts of the author(s) for the possible use of CIP Networks and do not reflect the ideas, opinions, and recommendation of ODVA per se. Because

CIP Networks may be applied in many diverse situations and in conjunction with products and systems from multiple vendors, the reader and those responsible for specifying CIP Networks must determine for themselves its suitability and the suitability of ideas, opinions, and recommendations expressed herein for intended use.

Copyright ©2007 Open DeviceNet Vendor Association, Inc. (ODVA). All rights reserved.
For permission to reproduce excerpts of this material, with appropriate attribution to the author(s), please contact ODVA on:

TEL	+1 734-975-8840
FAX	+1 734-922-0027
EMAIL	odva@odva.org
WEB	www.odva.org