NovaTech®
Bitronics ◆ D/3 ◆ Orion

D/3® Ethernet I/O Network Configuration

Introduction

With the release of the D/3® Ethernet Multi-Protocol Controller (EMPC2) interface in version 12.2-3, Ethernet I/O has quickly become one of the most popular D/3® interfaces. This interface is used to communicate with NovaTech's 8000 I/O and many common PLCs.

Ethernet as an I/O network has advantages such as the availability of software and hardware tools, networking hardware, and personnel who are familiar with the technology. Familiarity with Ethernet can also lead users to approach the design of the Ethernet I/O network as if they were designing an office Ethernet network.

The office Ethernet network that is used for email and internet access fulfills a different set of requirements than an Ethernet I/O network. Excessive network traffic on an office Ethernet network that delays a web page load or slows the spooling of a print job to a network printer will not impact production. Excessive network traffic on an Ethernet I/O network that causes delays can present safety concerns, impact production and damage equipment. The Ethernet I/O network requires determinism. This feature is achieved through the network design and configuration of the network hardware.

The typical office Ethernet network has no requirement for fault tolerance. The office network can sustain failures for periods of seconds and these failures would often be imperceptible to the users. In many industrial applications the Ethernet I/O network requires fault tolerance. The Ethernet I/O network used to interface to 8000 I/O provides industry leading fault tolerance at the network and I/O level.

The purpose of this White Paper is to explain how to properly configure an Ethernet I/O network for communication with the D/3®. Unlike the office Ethernet network, the Ethernet I/O network requires determinism and fault tolerance. These requirements are discussed and configurations are presented as examples that satisfy these requirements.

Ethernet Fundamentals

Ethernet is a CSMA/CD protocol.

• CS stands for carrier sense which is the equivalent of a telephone party line. Any node accessing the network is able to tell whether the network is currently being used.

• MA stands for multiple access. Any node may access the network. This is not so relevant for the D/3® Ethernet I/O interface since all supported protocols are query response. This means that nodes access the network in response to a query.

• CD stands for collision detection. When a collision occurs and is detected by the network interface card (NIC) the data transmission is ended temporarily. The NIC board waits for a random interval and then attempts retransmission. This process occurs until the data packet has been successfully transmitted. Collisions are a normal part of the operation of the network (Ethernet and TCP/IP based systems, E.J.Byres). However, excessive collisions can lead to unacceptable delays in data transmission and subsequent loss of Quality of Service (QOS). With a properly configured network where all the interconnected nodes are connected via switches, collisions are less of an issue since each switch port acts as its own bridged workgroup. This bridged workgroup connects to other bridged workgroups through the switching fabric of the switch. Any collision on a bridged workgroup is limited to that workgroup and does not affect the rest of the switched ports.

The design of our I/O network limits network communications as much as possible. The Ethernet I/O network is for communication between the EMPC2 interface and the I/O. The supported protocols are all query response; so that the EMPC2 queries the node and waits for a response. This alleviates some of the issues of media contention.

Collision of packets on the I/O network should rarely occur. If collisions occur with any regularity it is generally a sign that the network has not been configured properly. For example, late collisions may be a sign of a port duplex mismatch. Excessive collisions are an indication that there is traffic on the I/O network that should not be present. Ethernet I/O networks by design will not rise to the level of throughput that the I/O traffic results in large numbers of collisions.

Networking Considerations

Layer 3 Switch / Router

For simplicity, the layer 3 switch or router that is used to separate Ethernet I/O networks from the plant wide network will be referred to as a layer 3 switch. A properly configured router can be used just as well. The main difference between a layer 3 switch and a router is that the layer 3 switch has optimized hardware to do the actual switching. This often results in faster speeds. With 8000 I/O, a layer 3 switch or router is used to isolate the Ethernet I/O network from the plant wide network. The name layer 3 stems from the OSI communications model and the fact that a layer 3 switch interfaces at the Network layer of this model (the physical and data link layers are below this layer). The layer 3 switch acts as a router separating the traffic on the plant wide network from the traffic on the Ethernet I/O network. With the switch properly configured, only Ethernet traffic that should be on the Ethernet I/O network is allowed to pass, and this is limited to EBIM configuration traffic, asset management traffic, and diagnostic traffic. The layer 3 switch also isolates the Ethernet I/O traffic on one Ethernet I/O network from the other Ethernet I/O networks. The layer 3 switch provides multicast blocking and is configured for bootp passthrough.

Multicast Blocking

An 8000 I/O network must be isolated from plant traffic. Unnecessary plant traffic must be kept off the Ethernet I/O network and unnecessary I/O traffic must be kept off the plant network. There are two acceptable approaches. One approach is to limit access to the I/O network to the computer that runs the configuration software using a dedicated NIC card. This approach will work for a system that contains a single Ethernet I/O network. The second approach is to use a layer 3 switch to separate the I/O networks from each other and from the rest of the plant as shown in Figure 1. A properly configured layer 3 switch will prevent the 8000 I/O multicast traffic from getting passed onto the plant network and from getting passed to the other Ethernet I/O networks. The layer 3 switch will isolate the I/O networks from the plant traffic. This switch may be necessary so that personnel have access to the I/O network to access the HART information using a tool such as FieldCare. The switch in Figure 1 shows the incoming multicast and plant Ethernet traffic that is blocked by the layer 3 switch.

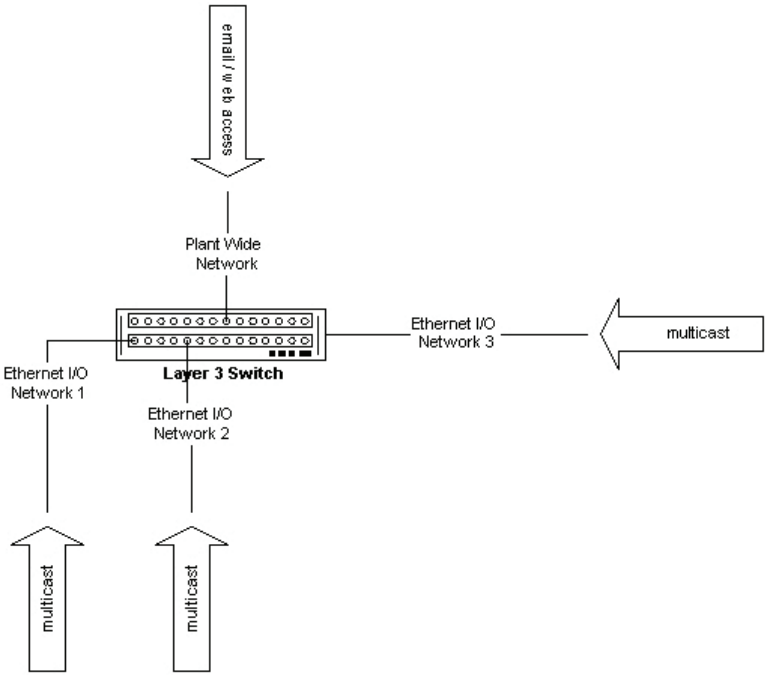

Figure 1 – Layer 3 Switch Configured to Isolate I/O Network

Bootp pass through

If the 8000 I/O network is isolated using a layer 3 switch, the switch must be configured for bootp passthrough. This ensures that the user is always able to access blank EBIMs on the I/O network.

Port Configuration

The port connecting the EBIM to the I/O network must be configured as autonegotiate. There is no way to hard specify the Ethernet ports on the EBIM.

The port connecting the EMPC2 to the I/O network must be configured as a hard set port. The port settings must match the settings in the D/3® Ethernet configuration tool. The EMPC2 port can be configured to 10 or 100 Mb, half or full duplex, the port settings on the switch must match. The EMPC2 cannot be configured as autonegotiate.

Switching Hardware

The switch hardware is the central element of the Ethernet I/O network. Office switches are not appropriate for the plant environment. Industrial Ethernet switches provide noise immunity, temperature performance and features such as dual power supply inputs. If the Ethernet network is in a controlled environment such as a control room, where noise and temperature considerations allow, office switches may be used. Wherever connections are being made through noisy environments in the plant, care must be taken to reject plant noise. For example, a fiber connection to a remote 8000 I/O drop could be used to integrate a remote Ethernet I/O drop into the control room PCMs. In many of the fielded systems, small switches are placed in the 8000 I/O cabinets, the EBIM connections are brought to these switches and a fiber connection is run back to the switch that connects to the PCM. Care must also be taken to use shielded twisted pair (STP) cabling in noisy environments for the short run cables.

Fielded systems and systems being shipped today are generally built with N-tron switches. A listing of some of the switches in fielded systems is provided in Table 1. The switch that connects to the EMPC2 must be managed so that the EMPC2 port can be configured to match the configured speed and duplex of the EMPC2 interface.

Manufacturer	Model	Type
N-tron	508tx-a	Industrial
N-tron	508fx2-a-st-s	Industrial
N-tron	9000 series	Industrial
Cisco	IE-3000-8TC	Industrial
Cisco	C2960	Office
Cisco	C3560G	Layer 3
HP	HP2626	Office

Table 1 – Fielded Ethernet Switches

Table 1 provides a sample of some of the switches fielded on D/3® Ethernet I/O networks. This is not a recommendation of specific switches. Details on standards compliance and additional hardware considerations can be found in "Industrial Ethernet on the Plant Floor", by Robert Lounsbury.

Network Fault Tolerance / Redundancy Approaches

One of the biggest differences between the office Ethernet network and the Ethernet I/O network centers on fault tolerance. A typical office Ethernet network will have limited or no fault tolerance. There are several reasons for this. The fault tolerance adds expense and complexity to the network and only provides a limited improvement in availability. The office Ethernet network changes on a day to day basis which makes the maintenance of any fault tolerance much more expensive. Finally, most office networking hardware is purchased over time, so any fault tolerance scheme would require the proper interaction of this feature across multiple hardware vendors.

Ethernet I/O networks on the other hand will often have some type of network redundancy and most often the recovery time from a network failure will be very fast.

Providing network redundancy is a difficult technical challenge. The most difficult part of redundancy in general is how you decide that a resource has failed. Unlike most redundant resources, networks are special because they are a shared resource. Assessing the network resource failure requires the cooperation of both ends of the network resource; for example the local and the remote node.

Because of this complication and the fact that this is a shared resource, there are two main approaches that are used to solve this problem. Professor Kirrmann of ABB differentiates these two approaches as static versus dynamic. Static and dynamic network fault tolerance are independent approaches to redundancy which must not be combined.

Dynamic Network Redundancy

Dynamic network redundancy is shown in Figure 2. With dynamic redundancy, the effective topology of the network changes to accommodate the failure. In Figure 2, the "x" marks a failed network connection. With the redundant connections available, the network uses an alternate path to the node. Note that the I/O nodes shown along the bottom of Figure 2 are singly connected to the network. The network redundancy therefore does not extend to the nodes since the single network connection presents a single point of failure. In the RSTP and self healing ring examples below, the network detects the failure and the network itself responds to the failure.

Rapid Spanning Tree Protocol (RSTP) and Self Healing Rings

RSTP and Self Healing Rings are dynamic approaches to network redundancy. RSTP and self healing rings solve the challenge of fault tolerance by investing the intelligence in the network. The network is designed with certain redundant paths. In order for the network to work properly with redundant paths, it is necessary that the redundant paths be removed or disabled until needed. Without removing redundant paths, nodes will receive the same packet several times and broadcast packets will circulate on the network forever. RSTP learns the network topology and then suppresses the transmission of packets on redundant paths. When the redundant path is needed, the RSTP is ready to reactivate these paths. Figure 2 shows a network configured with redundant paths. This network would require a protocol such as RSTP to remove these redundant paths from the active topology.

Self healing rings such as the Hirschmann Hyper Ring require a ring topology. This is a limitation that does not exist with RSTP. One switch acts as a ring master and breaks the ring (removes the one redundant path). The ring master is also responsible for sending packets around the ring to verify the continuity of the ring. When the ring master detects a break in the ring, the master closes the ring locally providing the redundant connection. This design allows the user to recover from the loss of a single connection in the ring. If the ring breaks in two locations, some degradation of communications would be expected.

A limitation of the RSTP and self healing ring approaches is that the knowledge of the network redundancy resides in the network equipment. In the case of a self healing ring, the ring master detects that the ring is closed or open by sending and receiving specific packets. These packets traverse the ring providing an indication that the ring is healthy. In the case of RSTP, the network switches "negotiate" among themselves to decide which redundant network paths need to be disabled. When RSTP detects a failure or change to the network RSTP activates

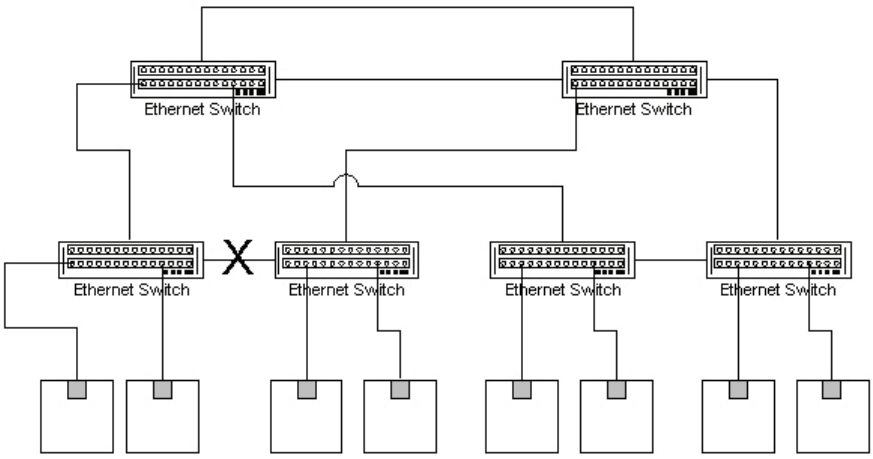

Figure 2 – Dynamic Network Redundancy[1]

[1] Hubert Kirrmann, Standard Redundancy Methods for Highly Available Automation Networks – rationales behind the upcoming IEC 62439 standard, ABB Switzerland, Corporate Research.

or deactivates the redundant path. In both self healing rings and RSTP configurations, the knowledge of network status resides within the network equipment and may not be immediately available to host applications.

A separate limitation of RSTP is that there is a learning process. This learning process occurs when equipment is powered on and when there is a change to the network topology. RSTP can recover from some topology changes in 1 second; however, during the learning process, RSTP sends many of the discovery packets around the network (Bridge Protocol Data Units). The Ethernet traffic from this learning process can be significant and can even rise to the level that this might appear like a broadcast storm to the 8000 I/O drops.

Static Network Redundancy

Static network redundancy is shown in Figure 3. With static redundancy, the topology of the network remains the same aside from the network or port failure, the node is responsible for selecting the port on which to communicate. In Figure 3, the "x" marks a failed network connection. The affected node detects the failure and switches active communications to the healthy port. Note that the I/O nodes along the bottom of Figure 3 are doubly connected to the network. With static network redundancy, redundancy extends to the I/O node. Static network redundancy places the network redundancy intelligence in the network nodes. This means the fault tolerant node must support this mode of redundancy.

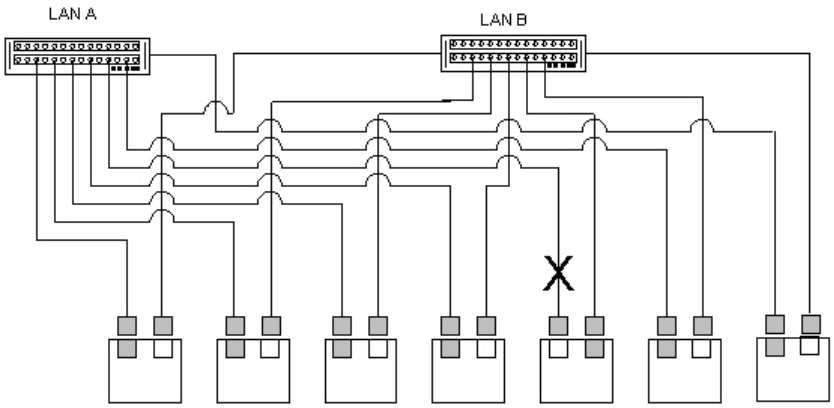

Figure 3 – Static Network Redundancy

Fault Tolerant Ethernet (FTE)

D/3® 8000 I/O uses Fault Tolerant Ethernet. FTE is a static approach to redundancy. All FTE nodes have redundant network connections. Each FTE node keeps track of network health and decides which port to use. The FTE network is designed so that the critical paths from the D/3® to the I/O are protected from single points of failure. FTE is designed to support FTE and non-FTE nodes.

FTE nodes use a multicast and broadcast scheme to determine network health. With non-FTE nodes, the FTE node sends out broadcast ARP packets to ensure connectivity. These ARP packets are sent on a configurable interval - every ARP timeout period which is defaulted to 10 seconds. With FTE nodes, the FTE node sends out a multicast message to ensure connectivity with all the FTE ports. These multicast packets are sent on a configurable interval – every pulse timeout period which is defaulted to every second. Based on who replies, the FTE node creates a map of network connectivity and therefore network health. The individual nodes on the network maintain information about the health of the network. Using this health information, the nodes are able to decide which Ethernet port to use in order to communicate with a remote node. In this redundancy scheme, the health of the network is maintained in the FTE nodes. The network itself acts as a resource and has no additional knowledge about network health or connectivity.

FTE Multicast

Multicast is an efficient means for one node to communicate with multiple nodes. In a multicast scenario, a multicast group is first established. The multicast group is established with a multicast query. This query asks nodes whether they want to join the multicast group. 8000 I/O EBIMs respond by joining the FTE group. The multicast scheme simply sends a multicast from each port of the FTE node, with the request that any FTE port respond. Based on the responses to this multicast, the FTE node builds a list of the ports that are visible from each port.

8000 I/O uses this network knowledge to provide information to the end user about the network status. If an 8000 I/O node is not able to see its own LAN B port from its LAN A port, the node deduces that connectivity between LAN A and LAN B is missing. This is a valid assumption for the local node to make since this connectivity is provided by the intra LAN link. The EBIM notifies the user that the intra LAN link has been lost.

Networking Determinism

For most of the applications that run on an office Ethernet network, there is no need for determinism. There is no need that the delivery of packets across the network be upper bounded in time. As the network experiences more collisions, more retransmissions occur and the average delivery time of packets will get longer. For an office Ethernet network, a considerable level of collisions can be easily tolerated. Excessive collisions ultimately indicate that the infrastructure is being overloaded. To quote E.J. Byres in Ethernet and TCP/IP Based Systems[2], "…Ethernet networks have been known to continue operation with collision rates as high as 40%". The Ethernet I/O network requires determinism, a bounded response time. In order to ensure a bounded response time, the Ethernet I/O

[2] E. J. Byres, Ethernet and TCP/IP-Based Systems, Instrument Engineers Handbook, Third Edition. ISBN 0-8493-1082-2.

Figure 4 – Non-Redundant PCM Non-Redundant EBIM Network Configuration

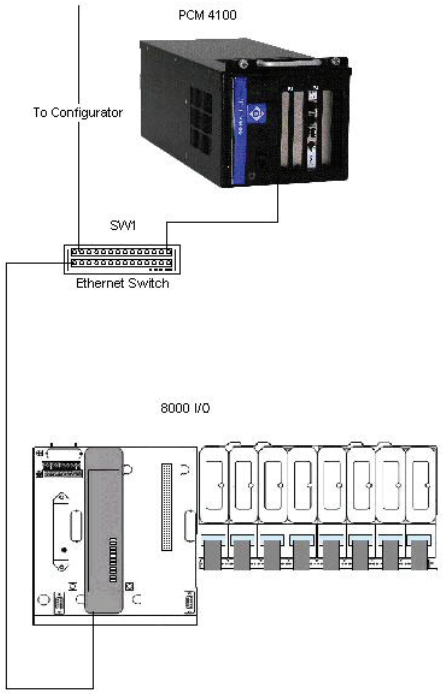

network traffic must be carefully controlled. To quote E. J. Byres[3] once again, "Ethernet delays are linear and can be consistently maintained below 2 ms for a lightly loaded network and 30 ms for a heavily loaded network".

The Ethernet I/O network must be isolated to ensure that the only Ethernet traffic on the network is I/O or configuration traffic. The I/O network can be isolated completely or partially. Connectivity with the I/O network is needed to configure the 8000 I/O drops and in some cases to gather diagnostic data. With a completely isolated I/O network, a separate NIC card can be inserted into the computer that will be used for 8000 I/O configuration. This interface can be connected to the I/O network. For a partially isolated network, a properly configured layer 3 switch provides connectivity to the I/O network while isolating the I/O network from plant traffic. Connectivity to the plant network may be needed for a couple of reasons. The software for configuring the EBIMs may reside on a remote computer that cannot be locally connected to the Ethernet I/O network. Access to the Ethernet I/O network may be needed for Plant Information systems such as FieldCare or any other asset management software the customer chooses to use. FieldCare accesses the HART® information by communicating directly with the 8000 I/O nodes. The general HART® data is accessible in the PCM if configured but is limited to the 4 HART variables (H1, H2, H3, and H4) with all the HART® status appropriately mapped. Finally, the Configuration software may need to be shared across several Ethernet I/O networks. This configuration requires that the traffic from one Ethernet I/O network be isolated from the other Ethernet I/O networks.

D/3® Ethernet I/O Network Configuration

Sample Configurations

The following provide recommended configurations for 8000 I/O Ethernet networks.

Non-redundant Configuration

- I/O network isolated
- All I/O and EMPC2 are on the same subnet
- Access to the I/O network for configuration is isolated by using a separate NIC for the CDCM

The configuration shown in Figure 4, provides the simplest connectivity. No redundancy is provided. A non-redundant PCM connects via a non-redundant network with non-redundant EBIMs. This network would look the same whether connected to 8000 I/O or to a Quantum PLC talking ModbusTCP. If the node is local, the PCM and the node connect into the same switch. If the node is remote, the remote node may connect to a remote switch that connects to the local switch with a shielded twisted pair cable or a fiber link. Depending on the distances involved and other considerations, the remote node may connect directly to the local switch via fiber or copper. If 8000 I/O is the node, only the A EBIM Ethernet port should be connected to avoid loading the EBIM with unnecessary broadcast traffic.

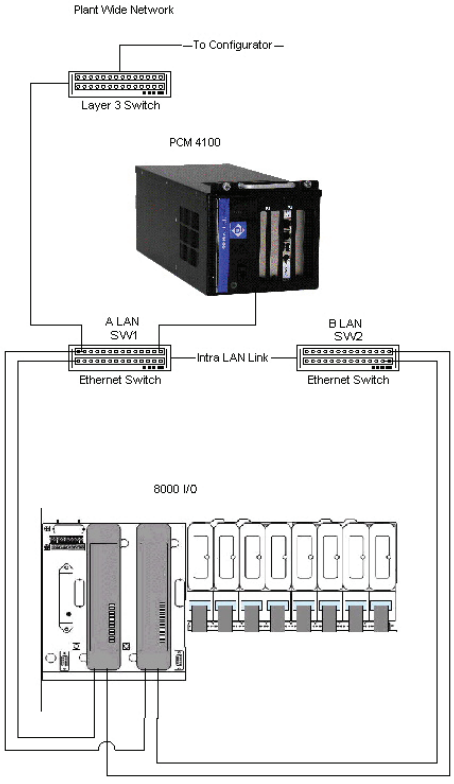

Figure 5 – Non-Redundant PCM Redundant EBIM Network Configuration

[3] D.R. Boggs, J.C. Mogul, C.A. Kent, Measured Capacity of an Ethernet: Myth and Reality, Proceedings of SIGCOMM '88 Symposium on Communications Architectures and Protocols, ACM SIGCOMM.

The switch port connecting the EMPC2 to the local switch is configured to match the EMPC2 setting in the Ethernet configuration tool. The switch port connecting the 8000 I/O node to the network is configured as autonegotiate. Pre-version 2.0 firmware only supported half duplex.

The I/O network is accessible to the configuration software via a third NIC.

Redundant EBIMs, non-redundant PCMs
- I/O Network isolated
- All I/O and EMPC2 are on the same subnet
- Access to the I/O network for configuration is isolated by a layer 3 switch
- Access to the redundant LAN relies on the intra LAN link

This configuration provides a level of redundancy that allows bump-less EBIM firmware downloads. Since this is a non-redundant PCM configuration, by definition there is no PCM redundancy.

The network configuration in Figure 5 (previous page) shows the non-redundant PCM EMPC2 interface connected to switch SW1. SW1 is a managed switch because the port the EMPC2 is connected to must be configured to match the port settings in the Ethernet configuration tool. SW1 and SW2 are connected by the intra LAN link. The single points of failure on this network are the EMPC2 interface and the connected switch SW1. Once the traffic has propagated to SW2 via the intra LAN link, the network is redundant out to the I/O rack.

The I/O network in this configuration is accessible to the CDCM via a layer 3 switch. This layer 3 switch allows the user to avoid the addition of a third NIC card in the CDCM. The CDCM is referred to as the "Configurator" since it is loaded with the 8000 I/O Configuration software. The layer 3 switch blocks the multicast traffic used by the FTE to determine network health and is carefully configured to allow bootp passthrough. This allows the CDCM to communicate with an uninitialized EBIM that is plugged into one of the racks. This only applies to the scenario where both redundant EBIMs are removed and an uninitialized EBIM is placed in the rack for programming. In most scenarios, the uninitialized EBIM will be placed in the rack next to a running EBIM and will be programmed by the running EBIM.

Redundant EBIMs, redundant PCMs
- I/O Network isolated
- All I/O and EMPC2 are on the same subnet
- Access to the I/O network is isolated by a layer 3 switch
- Access to the redundant LAN relies on the intra LAN link

In the most fault tolerant configuration, the user has redundant PCMs connected via a redundant network with redundant EBIMs. This configuration has redundant PCMs providing redundant operation and Online Upgrades. The redundant I/O network and redundant EBIMs provide redundant operations and online upgrades of firmware. The network will accommodate any single point failure (and many multiple point failures) notifying the operator of the failure and continuing to operate. Each of the switches, SW1, SW2, SW3, or SW4 can fail singly and the PCM will still communicate with the I/O. A redundant PCM / redundant EBIM network configuration is shown in Figure 6.

The FTE in the network in Figure 6 provides several layers of redundancy. Consider the following configuration:
- PCM01A selected
- RACK 1 EBIM A MASTER
- Communication path between PCM01A and RACK 1 via the EBIM A LAN A

Assuming a failure scenario where the A LAN connection to the EBIM is lost, the PCM will continue communicating with RACK 1 EBIM A, however the communication will pass through the EBIM A LAN B port. Connectivity will be as follows:
PCM01A <-> SW1 <-> intra LAN link <-> SW2 <-> EBIM A LAN B.

Assuming a follow on failure of the B LAN connection to RACK 1 EBIM A, the EBIMs will detect the network failures, and EBIM B will become MASTER. PCM01A will continue communicating with RACK 1.

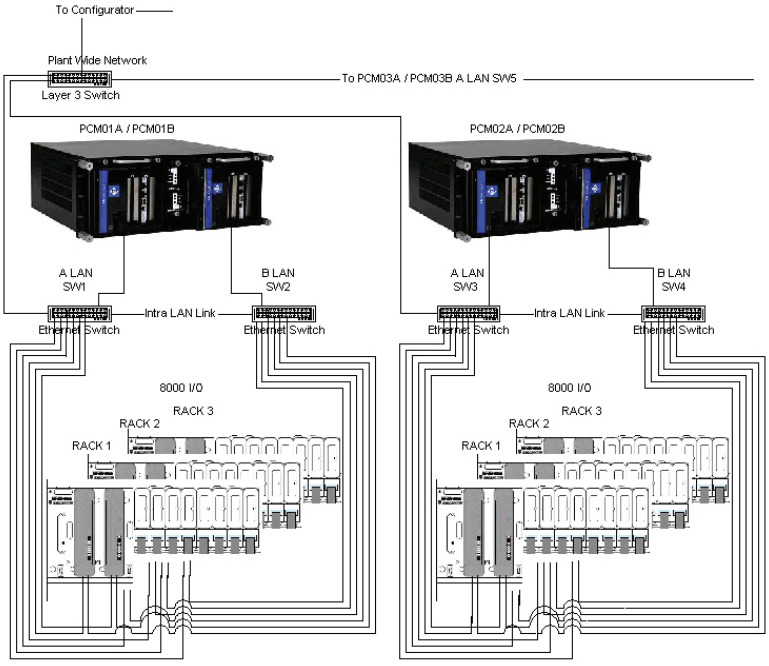

Figure 6 – Example Redundant Network Configuration

Assuming a failure scenario that involves the complete failure of switch SW1, if RACK 1 is configured as a failover node, the PCM01A will failover to PCM01B and communication will continue with RACK 1.

The failure scenarios make clear that the purpose of the intra LAN link is to make all ports on all EBIMs available to the EMPC2.

The failure scenarios also make clear that with a properly configured Ethernet I/O network, many network failures can be handled by the EBIMs either through switching which port is used for communication or by switching MASTERSHIP. Once these options are exhausted or in the scenario of certain network disruptions, the PCMs will fail over based on a failover node configuration.

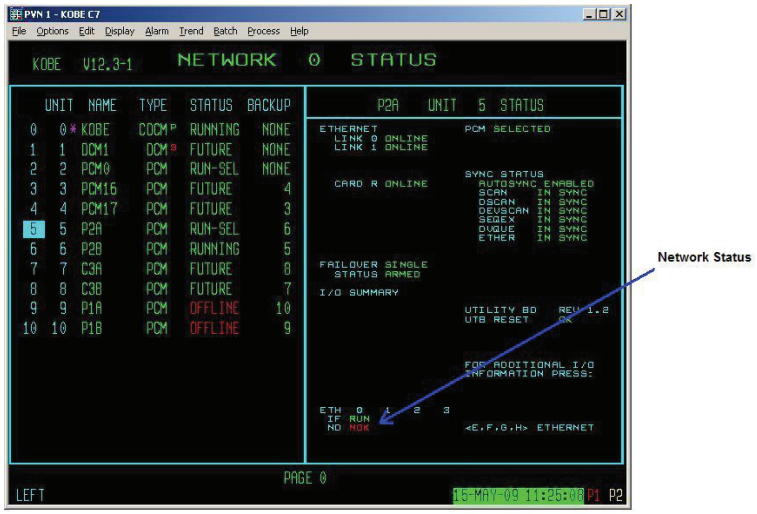

Figure 7 – System Status Display Overview

Redundant Configurations

The D/3® supports redundant Ethernet I/O networks. The configuration of these networks depends on the remote node being connected. Some of these redundant configurations are discussed next.

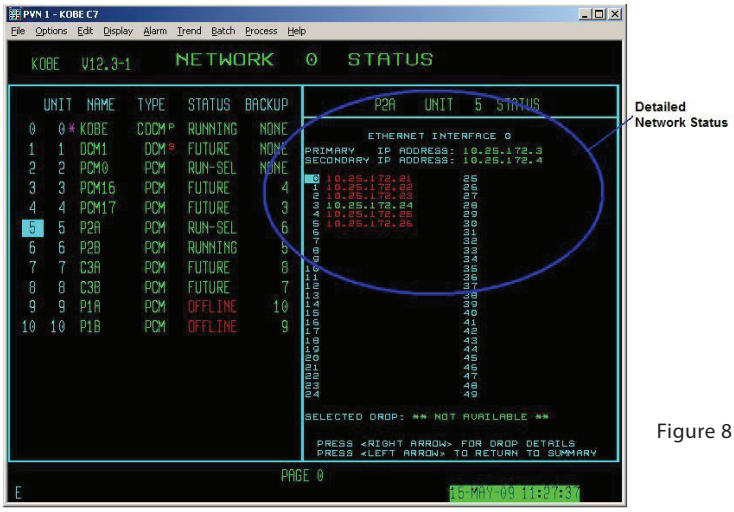

Figure 8 – System Status Display Network

Redundant Configurations

The D/3® supports redundant Ethernet I/O networks. The configuration of these networks depends on the remote node being connected. Some of these redundant configurations are discussed below.

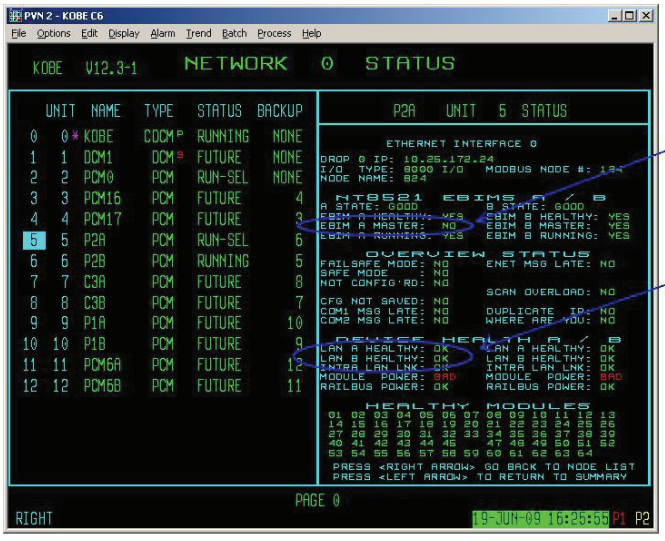

Figure 9 – System Status Display Node Details

8000 I/O

8000 I/O communicates with the D/3® using a modified ModbusTCP protocol. 8000 I/O uses FTE to provide network fault tolerance. Some of the benefits of this approach are discussed on the next page.

• I/O Status
The status of an EPN accessing 8000 I/O reflects all the possible failures that can occur from network level failure, to module level failure, to I/O point failure. For an EPN that experiences a loss of communication with the remote node, the EPN reflects this error with an IBAD/OBAD/BAD state. The user is able to view the network status using the system status display overview shown in Figure 7 and a more detailed display in Figure 8. For an EPN that experiences a module failure, the EPN reflects this error with an IBAD/OBAD/BAD state.

The user is able to view the module status on the system status network page for this node shown in Figure 9. For an AI/AO EPN that experiences a channel failure, the EPN reflects this error with the appropriately mapped hardware error UNDER/OVR/OCD.

• Network Status / Failover Status
The detailed system status shown in Figure 9 provides important status information on the EBIM pair and the network. The MASTERSHIP status is displayed. The status of each network connection and the status of the intra LAN link are displayed. This display provides the EBIM level view of the network and is one of the benefits of the FTE protocol.

• Direct A / Direct B Addressing
Each EBIM pair has 3 IP addresses, the MASTER, DIRECT A, and DIRECT B. All PCM communication with the EBIM pair passes through the MASTER address. This IP address can be shifted to any of the four EBIM Ethernet interfaces. Some of the engineering tools' access is through the DIRECT A and DIRECT B addresses. Care must be taken to make sure that all three addresses are accessible from the Configurator (machine running configuration software). These addresses must either be on the same subnet as the Configurator Ethernet interface or the Configurator Ethernet interface must be configured to route to that subnet.

• Broadcast Storms
The EBIM protects itself from broadcast storms by closing down its receive transceiver. This happens when the Ethernet messages in the Ethernet buffers meets or exceeds an arbitrary high water mark. When the transceiver is shut down, an event message is sent to the event log indicating "LAN Throttle". When the number of packets meets or is below the low water mark, the transceiver is turned back on. In the scenario where the EBIM throttles the LAN, all packets sent to the EBIM will be lost. If the LAN remains throttled for a long enough period, the EBIMs can lose communication with the PCMs. This would be indicated with a P3 alarm stating a node has changed from OK to FAIL.

• 1:6 Redundancy
The D/3® supports 8000 I/O redundancy that provides the user with a hot backup at the module level. This is called 1:6 Redundancy since the lowest level of redundancy provided is one spare for 6 active modules. This redundancy can automatically switch in a spare module to replace any of the active modules within 1 – 2 seconds of a failure. The module can be replaced with no disruption to the process and the user can perform a bump-less switch back to the replaced module. This allows the user to make full use of the spare modules, since they are ready to service the I/O at a moments notice. The user is also able to manually switch to the spare module for any maintenance that he might want to perform. For example, certain module firmware upgrades would be possible with no disruption to the I/O being serviced.

EtherNet/IP – ControlLogix
The EtherNet/IP Class 3 messaging protocol is a TCP/IP encapsulation of Rockwell's unscheduled Controlnet messaging. This is a query response protocol. The user requests tags from the ControlLogix controller and the ControlLogix controller responds.

• ControlLogix
ControlLogix controllers can be interfaced with the D/3® in both redundant and non-redundant configurations. The non-redundant configuration simply requires configuring the controller on the I/O network and interfacing with the controller via a 1756-ENBT module.

Rockwell supports redundant ControlLogix controllers by setting up identical ControlLogix chassis' and connecting them with a 1757-SRM module in each. The redundant information is passed across the 1757-SRM fiber link. The D/3® interfaces with these controllers via the 1756-ENBT

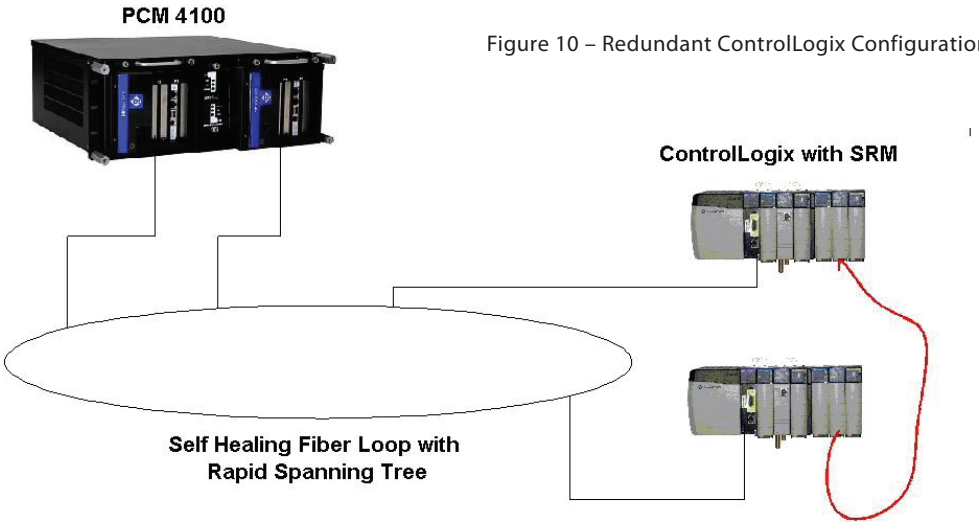

Figure 10 – Redundant ControlLogix Configuration

interface. Rockwell's redundancy approach is similar to the 8000 I/O. (The 8000 I/O redundancy however provides redundant Ethernet connections to each of the redundant EBIMs). The same IP address is shared between the redundant ControlLogix controllers and is only active on the controller that is "MASTER". This means that to the D/3® the redundant node appears as a single node. Network redundancy is provided independent of the redundant ControlLogix chassis through a self healing ring as shown in Figure 10 (previous page). As described above, the self healing ring will allow proper operation with the loss of a single connection in the ring.

ModbusTCP

• Siemens: No redundant Siemens configurations have been fielded. Siemens supports a ModbusTCP interface. The redundancy scheme for Siemens PLCs is different from other PLC vendors in the sense that there is an expectation that the Ethernet interface for each PLC resides actively on the same network. Therefore each PLC of the redundant pair has a unique IP address. The network redundancy would be provided by two self healing rings – one ring for each Ethernet interface. According to Siemens, it is possible to configure the two PLCs to the same IP address and isolate the Ethernet interfaces on separate networks. This allows a redundant PCM to interface with a redundant Siemens configuration. In addition, redundant Siemens PLCs do not care whether the user reads or writes to the Standby or Master PLC, the information is propagated to the other PLC.

• Quantum (non-redundant): Quantum PLCs have been fielded in non-redundant configurations.

• ABB (non-redundant): ABB PLCs have been fielded in non-redundant configurations.

• Triconex: Triple redundant PLCs have been fielded with the D/3® in a redundant configuration.

• Triplex (non-redundant): Triplex PLCs have been fielded with the D/3® in a non-redundant configuration.

• Gateways (non-redundant): Gateways provide access to data that might be physically impossible to bring directly into the PCM. In some cases, the data is remote from the PCM and there is no other method of bringing the data back.

o Orion: The Orion supports many serial protocols. The Orion interfaces to the D/3® via ModbusTCP. Using the Orion, it is possible to "concentrate" a large number of serial devices for communication with the D/3®. The Intelligent Electronic Device (IED) protocols supported by the Orion 5r are as follows:

ABB DPU
Allen Bradley DF1
Areva KITZ
Areva Optimho
Basler DFPR
DNP3 Serial and IP
GE DLP
GE Moisture Meter
GridSense PAC
IEC 870-5-103
Keithley Meter
Modbus Serial and TCP
PG&E 2179
RFL
SEL® ASCII
SEL® Fast Meter
SEL® Fast Operate
SEL® Fast SER
SPA Bus
TransData DTO

With the Orion configured with several serial ports, it is possible to concentrate several serial networks into one node that can communicate with the D/3® via ModbusTCP.

o Anybus AS-Interface: AnyBus supports protocols such as AS-interface. These gateways are accessible via ModbusTCP.

o ModbusTCP – serial modbus: ModbusTCP gateways allow the user to communicate with remote serial modbus devices anywhere in the plant.

Multiple protocols on the same network

It is recommended that 8000 I/O be isolated on its own Ethernet interface and Ethernet I/O network. This is based on several factors including the specific network requirements of FTE and 8000 I/O. EtherNet/IP Class 3 messaging is also recommended as an isolated interface and network. ModbusTCP and DF1 over Ethernet can be configured on the same network and interface.

All protocols have been tested on the same interface at the same time. With a lightly loaded system, more can be done with a single interface. The variable that needs to be considered is the number of blocks that is being supported on the interface. Since the Ethernet I/O network is a shared resource consideration must be made to the loading of a network when multiple protocols are being supported on a single interface. What is theoretically possible is not always the best implementation.

Summary

The typical office Ethernet network has no requirement for determinism and may have some limited redundancy. The web page that does not load on the initial request is annoying but functionally insignificant. Network components that fail are quickly detected and fixed by the IT department. The Ethernet I/O network must be designed and implemented to provide deterministic and fault tolerant communications with I/O. These are very different requirements than the office Ethernet network. The appropriate configuration and equipment must be considered for this purpose.

The EMPC2 supports multiple protocols making it possible to interface with many different PLCs. Several of these PLCs support redundant configurations and redundant networks. Most PLCs rely on RSTP or a self healing ring to provide network redundancy.

The EMPC2 interface provides deterministic and fault tolerant communications with 8000 I/O. The proper design and implementation of the Ethernet I/O network is essential for the proper operation of the FTE protocol. The FTE protocol provides a higher degree of redundancy than either RSTP or a Self Healing Ring. The FTE protocol integrates the network redundancy into the system in a way that allows the user to monitor the Ethernet I/O network health.

With a properly designed and configured Ethernet I/O network, the D/3® is a proven solution for deterministic and fault tolerant communication with 8000 I/O.

Contact:

NovaTech, LLC
D/3 Division
11500 Cronridge Drive, Ste. 110
Owings Mills, MD 21117

T: 410.753.8300
F: 410.753.8395
E: info@novatechps.com
www.novatechweb.com