

the **EXTENSION**

A Technical Supplement to Control Network

© 2005 Contemporary Control Systems, Inc.

Power Over Ethernet—Industrial Ethernet is Now a True Fieldbus

By George Thomas, Contemporary Controls

Industrial Ethernet as a Fieldbus

Industrial Ethernet, as a fieldbus replacement technology, lacks two attributes found in fieldbuses. The first is bus topology. The second is power sourced from the network cable to energize field devices. DeviceNet is a good example of a fieldbus that can accomplish both. DeviceNet can be wired in a bus topology while providing 24-volt power in the cable for powering field devices such as photo-eyes, push-button stations, and limit switches. Higher-powered actuators usually have their own power sources. Industrial Ethernet only supports star topology and, until now, could not provide power over the cable without implementing a non-standard approach. With the approval of IEEE 802.3af in 2003, the power-sourcing problem has been solved with the Power Over Ethernet (PoE) standard. PoE not only provides a safe and effective way of applying power to modern devices, but also utilizes its star topology to its advantage by controlling the amount of power each connected device receives while protecting non-powered devices from harm. DeviceNet is incapable of doing the same with its bus topology.

Twisted-pair Cable Carries the Power and Data

PoE was not designed for the industrial or building automation markets but for a much larger information technology market incorporating IP phones, IP cameras, and wireless access points. An IP phone, which utilizes Voice over IP (VoIP) standards, should look and feel like an ordinary telephone. A telephone is powered from its data connection and works when the power is lost because its 48-volt power source operates from batteries. These same attributes are attractive in industrial and building automation when it is inconvenient or expensive to run higher voltagepower in the field or when it is desirable to backup the entire automation system from one power source. Most automation systems operate from 24 volts and not 48 volts so the 802.3af standard must be examined for applicability to the automation industry.

Modern Ethernet cabling, complying with the 10 Mbps 10BASE-T and 100 Mbps 100BASE-TX standards, consists of four twisted-pairs of which only two pairs are employed. The Gigabit standard, 100BASE-T, uses all four pairs. Pre-802.3af compliant networks attempted to utilize the unused pairs for carrying power by simply applying 24 or 48 volts to the unused pairs assuming they were not being employed for any other purpose. Clearly this would not work for Gigabit Ethernet since all pairs are used for data. The other approach is to incorporate the data lines for carrying power. This would allow Gigabit devices to be powered as well. The 802.3af standard provided a unified approach to powering devices using either method while addressing the issue of not damaging non-PoE compliant devices that are not expecting power. The standard is much more extensive than simply applying power to unused pairs, as we shall find out.

PSE and PD

The 802.3af standard identifies two types of devices. The Power Sourcing Equipment (PSE) provides the required power just like its name implies. The Powered Device (PD) is the receiver of the power. The two are connected via the Ethernet communications cable. Understanding Ethernet's star topology, the PSE would naturally be assumed to be a port on a hub or switch while the PD would be an end station or node on the network. There are two ways of connecting power as shown in Figures 1 and 2. The first method, Alternative A, is via the data lines while the second method, Alternative B, is via the unused pairs. Using both approaches simultaneously is currently not allowed. The standard is being reviewed to determine if by using both approaches, more power could be transferred to the PD. When using the unused pairs, the wires within the pair are connected together in order to increase the cable's current carrying capacity. The same trick can be used with the data lines, but in this case the power must be fed through a center-tapped transformer. If the DC-current flow in each wire of the pair is equal, there will be no DC bias established in the winding of the transformer so this method could be very

effective for carrying power and data simultaneously. In fact, this method is the only one possible for using Gigabit Ethernet since no spare pairs exist. A PSE can be designed to source power either through the data lines or through the unused pairs. The PD must be able to do either, making the PD design a bit more complex.




Figure 1—Alternative A. The signal pairs carry both data and power. Polarity is indeterminate.





Figure 2—Alternative B. The unused pairs carry power. Polarity is as shown.

The power being supplied by the PSE is a nominal 48 volts (44 to 57 volts). When using the spare pairs for power, the voltage polarity is defined. When using the data pairs for power, the voltage polarity is undefined requiring the PD to be able to accept either polarity and still function. The reason being that a crossover cable could be between the PSE and the PD. Normally, a straight-through cable is used to connect a switch to an end device (see Figure 1). The end device is wired as an MDI while the switch is wired as an MDIX. The “X” means that the switch implements an internal crossover function by having transmitters on the switch connect to receivers on the end device, and receivers on the switch connect to transmitters on the end device. If two MDI-compliant devices need to communicate to one another (an end device to another end device), then a crossover cable is required to make the equivalent connection. Modern switches have a feature called Auto-MDIX where the switch port will adjust to either a straight-through cable or a crossover cable and still function with an end device or another switch port. A crossover cable matches up data pairs by having pins 1 and 2 on one side connect to 3 and 6 on the other. Likewise,

pins 3 and 6 are matched up with 1 and 2. Because of this crossover situation, the polarity of the voltage on the data pairs will reverse at the PD. To guard against this reverse voltage condition, an auto-polarity circuit should be used at the PD. Crossover cables do not affect the spare pairs. Therefore, the PD will experience no polarity change. To be safe, it is best to incorporate auto-polarity on the spare pairs as well.


End-point and Mid-span PSEs

There are two kinds of PSEs. The first is an end-point PSE (Figures 3 and 4) where the power sourcing equipment is within the hub or switch. The Ethernet switch now becomes a Power Sourcing Ethernet switch adding complexity and expense. However, field cabling is not disturbed. The standard Ethernet switch is now replaced with a PoE Ethernet switch. In terms of data handling, the PoE switch operates identically to that of a non-PSE switch. Additional circuitry and an adequate power supply are necessary to serve the powered Ethernet ports, but not all ports need to be powered and frequently this is done to reduce power supply requirements. A switch-to-switch cascading connection (uplink) would not need powered ports so non-PoE ports would be present in PoE switches.



Endpoint PSE, Alternative A

Figure 3—An end-point PSE can utilize Alternative A or B. Alternative A is shown above.

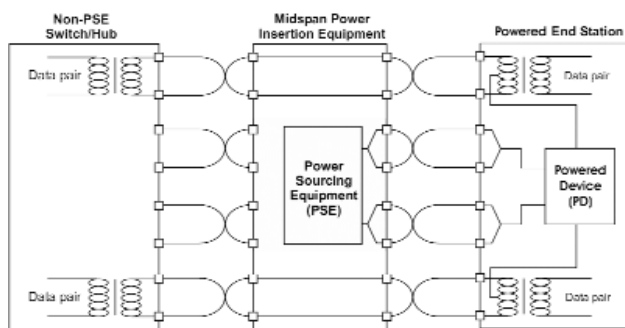


Endpoint PSE, Alternative B

Figure 4—Alternative B utilizes the spare pairs for carrying power. No specialized data transformers required at the PSE.

The second PSE is the mid-span PSE (Figure 5) where the mid-span device can reside anywhere within

the 100-meter length of Ethernet cable attaching a conventional hub or switch port to an end station. The advantage of this device is that standard off-the-shelf Ethernet switches can be retained in PoE applications since the power comes from the mid-span device and not the switch. In this application, the Ethernet switch is not disturbed, as is the field cabling. The mid-span device is not a hub or a switch but mostly resembles that of a patch panel with a conventional Ethernet input port and a corresponding powered output port. The downside of it is that it doubles the connections and introduces another piece of equipment to mount and energize. Mid-span devices are beneficial for testing the PoE concept or for operating with legacy equipment during a retrofit.



Midspan PSE, Alternative B

Figure 5—A mid-span device can only use Alternative B and, therefore, will not operate with Gigabit Ethernet.

Alternatives A or B

As mentioned previously, there are two ways of energizing end devices. Power can be applied to either the data pairs or the spare pairs but not both. An end-point PSE vendor makes the choice. On the other hand, the PD vendor has no choice and must be able to handle either method. Notice in Figure 4 that the PSE vendor decided to use the spare pairs (Alternative B) to power the PD. In this case, no special consideration was given to the transformers on the data lines. Since the PD must accommodate either Alternative A or B, the PD must utilize center-tapped transformers. The PSE vendor could have easily decided to use the data pairs (Alternative A) as shown in Figure 3. The vendor is allowed to support both approaches but not at the same time.

In Figure 4, the mid-span PSE vendor utilized Alternative B. This is the only choice for a mid-span device since the data lines must be passed without disruption. Since spare pairs are being used, a mid-span device cannot be used with Gigabit Ethernet. Only Alternative A systems can support Gigabit Ethernet. In this situation, the end-point PSE and PD would both need four, center-tapped transformers.

Regardless of which alternative is being used, there is an electrical isolation issue for the PD. The standard specifies 1500 volt RMS isolation. Not only must the transformers be so rated, the DC-DC converter within the PD must have the same isolation. This eliminates the lower cost converters that have no input transformer. A proper implementation would require transformer isolation on the input to the DC-DC converter. Again, this increases the cost of the PD.

Device Identification and Classification

Power management is possible with 802.3af standard. In order to protect non-PoE compliant devices from receiving unexpected power, an identification process is performed during connection time. Un-attached PSE ports are un-powered in their dormant state. When a device is first attached to a PSE port, it must present itself as a nominal 25k ohm resistance. This is called its Detection Signature, which the PSE tries to identify by applying up to 10.1 volts while measuring the current. Once identified, the PD may or may not participate in a classification process, which is used to determine the power that must be reserved for the PD. There are currently four classes of power (0 through 3) and one reserved class(class 4). Class 0 is the default class indicating that the PD may use a minimum of 0.44 watts to the maximum allowed, which is 12.95 watts. With the default setting, there is no power management. Notice the minimum power requirement. The PSE will remove from devices that draw less than the minimum required. Classes 1 through 3 offer a range of power as shown in Table 1. The PSE will apply a higher voltage (up to 20.5 volts) in order to determine the class of device by measuring the resulting current. This current signifies the class and is called the Classification Signature. If none of the anticipated currents are measured, the device is deemed to be a class 0 device and no power management is possible. In either case, the output voltage is then increased to its nominal 48 volts while its output current is monitored.

PD Power Classification

Class	Usage	Range of maximum Power Power used by the PD
0	Default	0.44 to 12.95 Watts
1	Optional	0.44 to 3.84 Watts
2	Optional	3.84 to 6.49 Watts
3	Optional	6.49 to 12.95 Watts
4	Not Allowed	Reserved for Future Use

Table 1—Four power classes are currently available, however, only three classes have the possibility of being managed.

Injectors and Taps

Some PoE schemes were developed before the 802.3af standard was ratified and are still in use. These schemes may or may not be compliant to the standard. They utilized the spare pairs to power legacy devices, however, the concept is still applicable today. Usually an injector is used to apply power to an Ethernet segment, and a tap is used to extract the power from the segment before it appears at a non-PoE end station. To follow the standard, the injector must apply a nominal 48 volts, however, products exist that provide non-compliant voltages (such as 24 volts). There are two types of taps. The passive tap simply removes the power before the end station and then feeds it directly to the power input of the legacy device. The regulated tap will adjust the voltage on the cable to exactly match the requirements of the legacy device. This approach can be very effective with single station applications, however, to be truly compliant to the 802.3af standard, the tap must participate in the signature detection process and the injector must comply with all the requirements of a mid-span PSE. This may not be the case.

Power and Network Management

Each port on a PSE device must be capable of delivering 350 mA at 44 volts. This yields 15.4 watts of power, which is the maximum for a port. However, the PD cannot draw more than 12.95 watts. Power drops are expected with 100 m of twisted-pair cable accounting for the difference. For most automation applications, 13 watts of power is quite adequate. Although each port can handle 15.4 watts, it does not mean that the PSE needs to be rated for the maximum at each port. This is where power management comes in. By having PDs report only the power necessary, it is possible to undersize the power supply and reduce its expense for the number of ports that are supported.

Another possibility is to equip the PSE with the Simple Network Management Protocol (SNMP). By introducing Management Information Base (MIB) data for each of the powered ports, power cannot only be monitored but also controlled. This information may then be used for security purposes or remote troubleshooting. Powered end-stations can then be rebooted remotely by the removal and application of power.

Summary

Power over Ethernet involves more issues than simply defining the cable connections, and the 802.3af effort provides valuable guidance on how it should be accomplished. Although many experts feel that device-level Ethernet is still too costly compared to equivalent fieldbus devices, POE makes a significant step towards Industrial Ethernet as a fieldbus replacement.

REFERENCES

IEEE 802.3af-2003 Amendment to IEEE 802.3-2002.
Poweroverethernet.com.

CONTEMPORARY CONTROLS®

www.ccontrols.com

www.ctrlink.com

www.arcontrol.com

Past issues of the copyrighted Extension are available. Please visit our web site www.ccontrols.com. Select Support and click on Extension Archive.