

The ABCs of Power Over Ethernet

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 recently, 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 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 Carries 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 voltage power 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, 1000BASE-T, uses all four pairs. Prior to 802.3af, some networks utilized 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. Another approach was to use the data lines for carrying power. This would allow Gigabit devices to be powered as well. The 802.3af standard provides a unified approach to powering devices with 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 learn.

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) receives the power. The two are connected via the Ethernet communications cable. Given 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 illustrated in Figures 1 and 2. The first method, Alternative A, uses the data lines while the second method, Alternative B, uses the spare 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 spare pairs, the

wires within the pair are connected together to increase the current carrying capacity. The same capacity can be obtained with the data lines if the power is 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 with Gigabit Ethernet since no spare pairs exist. Different PSEs can be designed to source power differently, one using the data lines and the other using the spare pairs. One PD must be able to work with either type of PSE, making the PD design a bit more complex.




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




Figure 2 — Alternative B. Spare pairs carry power. Polarity is indicated.

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 polarity is undefined. This requires the PD to accept either polarity and still function — because a crossover cable might connect the PSE and the PD. Normally, a straight-through cable is used to connect a switch to an end device (as in Figure 1). The end device is wired as an MDI and the switch is wired as an MDIX. The “X” means that the switch implements an internal crossover function whereby 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 for the equivalent connection. Modern switches have a feature called Auto-MDIX where the switch port adjusts to either a straight-through or crossover cable and still function with an end device or another switch port. A crossover cable matches data pairs so that 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. Consequently, the voltage polarity on the data pairs will invert through a cross-over cable. To guard against this reverse-voltage condition, an auto-polarity circuit must be used at the PD so the PD will experience no polarity change. Crossover cables do not affect the spare pairs, but 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. In terms of data handling, the PoE switch is identical to 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 often this is done to reduce power supply requirements. A switch-to-switch cascading connection (uplink) needs no powered ports, so non-PoE ports could be present in PoE switches.




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





Figure 4 — Alternative B uses the spare pairs for power, so no special transformers are required.

The second type of PSE is the *mid-span* PSE (see Figure 5) which can reside anywhere within the 100 m 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, neither the switch nor the field cabling is disturbed. The mid-span device is not a hub nor a switch but mostly resembles 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, so it can not work 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 power the PD via the spare pairs (Alternative B). No special transformers were used on the PSE data lines. The PSE vendor could have used the data pairs (Alternative A) as shown in Figure 3. Since the PD must accommodate either Alternative A or B, the PD must use center-tapped transformers. The vendor can support both methods but not in the same piece of equipment.

In Figure 5, the mid-span PSE vendor utilized Alternative B — the only choice since the data lines

must pass through the mid-span without disruption. Since spare pairs are 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 each 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 for the transformers and the DC-DC converter in the PD. This eliminates lower-cost converters that have no input transformer. Proper implementation requires 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. To protect non-PoE compliant devices from receiving 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 for determining the power that must be reserved for the PD. There are currently four classes of power (0–3) and one reserved class (class 4). Class 0 is the default class indicating the PD will use a minimum of 0.44 watts or as much as the maximum allowed 12.95 watts. The default setting uses no power management. Notice the minimum power requirement. The PSE will remove power from devices that draw less than the minimum. Classes 1–3 offer a range of power (see Table 1). The PSE will apply up to 20.5 volts 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 class 0 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 Power Used by 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, but only three may be managed.

Injectors and Taps

Some PoE schemes were developed before the 802.3af standard was ratified and are still in use. These schemes, that may or may not be compliant to the standard. They utilized the spare pairs to power legacy devices — a concept 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 gets to a non-PoE end station. To follow the standard, the injector must apply a nominal 48 volts. 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 stations, but to comply 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 with up to 100 m of twisted-pair cable account 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 powered port, power cannot only be monitored but also controlled. This information may then be used for remote troubleshooting or security purposes. 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 toward Industrial Ethernet as a fieldbus replacement.

REFERENCES

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

Poweroverethernet.com.

