New Power for Ethernet—Powered Devices (Part 2 of a 3-Part Series)

by Dave Dwelley

> **Powered Ethernet Promises to Remove Warts**
>
> *An excerpt from Part 1 of this series:*
> For years, data has passed over Ethernet CAT-5 networks, primarily to and from servers and workstations. The IEEE 802.3 group, the originator of the Ethernet standard, is currently at work on an extension to the standard, known as 802.3af, which will allow DC power to be delivered simultaneously over the same wires. This promises a whole new class of Ethernet devices, including IP telephones, wireless access points, and PDA charging stations, which do not require additional AC wiring or external power transformers ("wall warts").

Introduction

An IEEE 802.3af Powered Ethernet connection provides both the familiar 10/100/1000MB/s data link and 13W worth of 48V DC power to a connected device. Such a device, known as a PD (for Powered Device), can be a digital Voice-Over-IP phone, a network wireless access point, a PDA charging station, an HVAC thermostat, or almost any small Ethernet-connected data device that would otherwise be powered by a wall transformer. A PD need not use the data link at all; something as simple as a cell phone battery charger or an illuminated exit sign could draw its power from an Ethernet connection.

This article is the second in a three-part series on Powered Ethernet. This issue covers the operation of the PD in detail. Part 1 appeared in the last issue of the *Linear Technology* magazine and covered the power details of the system, with a focus on the PSE (Power Sourcing Equipment) and its characteristics. Part 3 will discuss the nuances of detection and classification—the mechanism that the 802.3af standard uses to ensure that PDs receive power while legacy data-only devices remain unpowered.

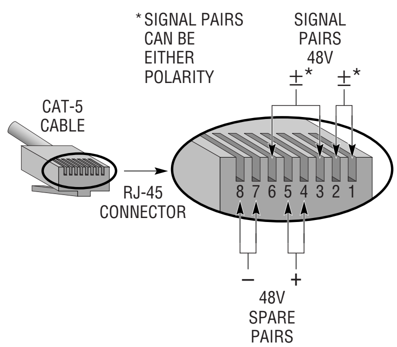

Figure 1. Signal and spare pairs on RJ-45 connector

Characteristics of a PD

Power arrives at the PD on a standard CAT-5 network cable via an RJ-45 connector. A CAT-5 wire contains four twisted pairs of 24-gauge wire (8 conductors in total). Two of the pairs—the "signal" pairs, at pin pairs (1, 2) and (3, 6) on the RJ-45 connector, shown in Figure 1—are used for the standard 10/100 Ethernet transmit and receive links. The two other pairs—the "spare" pairs, at pin pairs (4, 5) and (7, 8) are unused in 10/100 networks. 1000BASE-T networks use all four pairs. 48V appears on the cable as a difference in the common-mode voltages between the two signal pairs or the two spare pairs (but never both). The signal pairs are transformer-isolated as they enter the PD to strip the DC out of the data signal path; the power is taken from center taps on these transformers and passed to the PD input circuitry, as shown in Figure 2. The spare pairs may or may not be transformer isolated.

To be considered an IEEE-802.3af PD, a device must meet several criteria. A PD must be able to accept power over either the signal pairs or the spare pairs, since a PSE is allowed to power either set. This is typically accomplished by diode ORing the two power inputs, as shown in Figure 3a. This circuit has the additional advantage of removing the signature from the unused set of pairs when the power is applied to the other set, a requirement of the IEEE spec. PDs are not allowed to draw power from both sets of pairs simultaneously. Diode bridges can be used to implement auto-polarity; this is useful since many CAT-5 cables are wired as crossover cables, so voltage polarity is likely to arrive reversed instead of forward. An alternate connection, shown in Figure 3b, uses single diodes and a third reverse biased diode to present an invalid signature when the polarity is reversed. This circuit will work with 802.3af systems, although it will not be powered if a crossover cable is used.

All PDs must present a characteristic 25kΩ signature impedance at the power inputs when probed with voltages between 2.8V and 10V. The signature impedance is allowed to

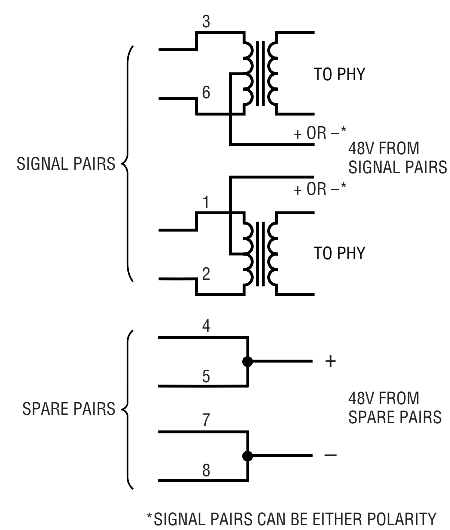

Figure 2. Deriving power from the cable

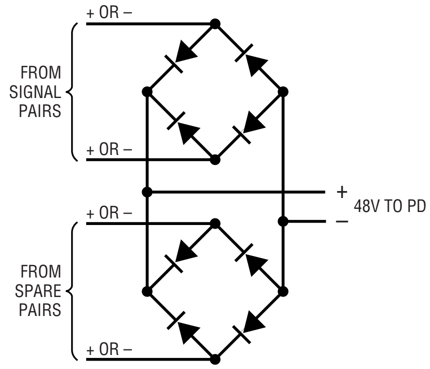

Figure 3a. Autopolarity input circuit

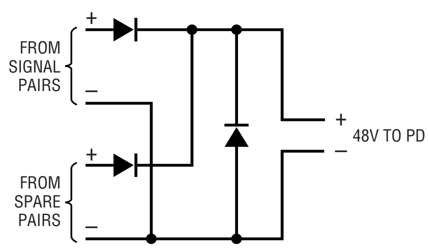

Figure 3b. Non-autopolarity input circuit

Table 1. PD power classifications and signature currents

Class	PD Maximum Power	Nominal Classification Signature Current
0	0.44W–12.95W	< 5mA
1	0.44W–3.84W	10.5mA
2	3.84W–6.49W	18.5mA
3	6.49W–12.95W	28mA
4	Class 4 is currently reserved and should not be used	40mA

have up to three diodes in series, to allow for diode-based power steering and autopolarity circuits. This signature is an indication to the PSE, typically the Ethernet switch or hub, that the device on the end of the wire is, in fact, a PD, and won't be damaged if the PSE applies 48V to it. Older Ethernet devices, such as network interface cards and non-powered hubs, typically present common-mode impedances of around 150Ω, well away from the valid PD impedance.

A second, optional signature may be presented at the terminals when probed with between 15V and 20V. This "classification" signature indicates to the PSE the maximum power the PD will draw so the PSE can budget power if it chooses to. The classification signature appears as a constant current drawn by the PD at the input terminals. Table 1 shows the classes and their constant current signatures. Classes 1, 2 and 3 are used when the power is known. Class 0 is assigned if the PD chooses not to implement the classification signature. Class 0 means the PSE does not know how much power the PD may draw, although it's generally wise to budget Class 3 power for such a PD. Class 4 is reserved for future use.

Once the PD has identified itself to the PSE, the PSE will apply a voltage between 44V and 57V to the wire. The PD now has several obligations. It should not draw significant load current until the terminal voltage rises above 30V (to avoid interfering with the classification signature), yet it must be fully operational by the time the line voltage reaches 42V. It can never draw more than 350mA or 12.95W continuously, whichever is less (brief surges to 400mA are allowed under some circumstances). It needs to operate with as much as 20Ω of wire in series with the input, which can cut the input voltage by as much as 8V during a 400mA current surge. This mandates adequate hysteresis between the turn-on and turn-off voltages to prevent *motorboating*—oscillating on and off—when the load is first applied and the input voltage is low. The PD must have an input capacitance below 180µF to keep the power-on current surge to a reasonable level; if this input capacitance is larger than 180µF, the PD must actively limit the inrush current to keep it under 350mA. Finally, the PD must maintain at least 10mA of current draw and must maintain an AC impedance of 33kΩ or less to avoid being disconnected.

LTC4257 PD Power Interface Controller

The LTC4257, shown in Figure 4, is designed to satisfy the specific demands that the IEEE standard places on a PD, allowing designers to focus on overall system design without worrying about compliance. The LTC4257 includes a trimmed 25k signature resistance on-board, and a full classification signature circuit, programmable to classes 0, 1, 2, 3 or 4 with a single external resistor. An on-chip power MOSFET keeps the PD circuitry disconnected from the line until the voltage rises above 40V. An inrush current limiting circuit keeps the line current below 400mA at all times, and thermal limiting protects the circuit from extreme fault conditions. The only task passed on to the rest of the PD circuitry is keeping the continuous power drain under 12.95W (or lower, if Class 1 or 2), something a switching regulator, such as the LT1871, does automatically. The regulator circuit must also maintain the required minimum 10mA current draw, a requirement usually met by the quiescent operating current of the system. *continued on page 27*

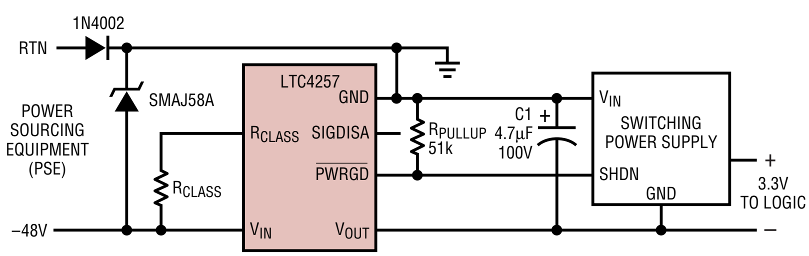

Figure 4. Typical PD application

LTC4257, continued from page 10

Two additional features add flexibility to LTC4257 designs. An open-drain $\overline{PWRGD}$ output indicates that the voltage drop across the internal power MOSFET has dropped below 1.5V, indicating that any input capacitance has charged, the output has reached its final value, and it is safe to turn on the system. This helps systems that draw the maximum input power stay below the inrush limits at turn on. A SIG_DISA input allows the PD to disable the 25k signature resistance if desired, allowing it to opt not to receive power from the PSE if it is getting it from another source, such as a wall transformer.

Conclusion

The LTC4257 contains virtually all of the circuitry needed to connect a powered device to an IEEE 802.3af Power Over Ethernet network. Signature, classification, power switching, inrush, and fault protection are all included, thus simplifying the required circuitry between the input transformers and the PD voltage regulator. The LTC4257 accomplishes all of this in a space-saving 8-pin SO or DFN package with only one external component, a resistor to program the class current (not needed for class 0).

Part 3 of this series will cover the details of detection and classification from the PSE end of the power network.