

New Ethernet systems distribute DC power with data

The universal network is coming to an RJ-45 jack near you, carrying power with your packets!

Power-over-Ethernet (PoE) is a power-distribution technique recently approved as IEEE standard 802.3af. It allows that ubiquitous and universal network called the Ethernet to carry DC power along with your data packets. The recent ratification of this IEEE standard appears to be the first instance in which a universal power jack or plug has been defined for worldwide use. It means that local AC power is no longer needed for all network-attached devices that need continuous power—such as IP phones, wireless access nodes, and web-surveillance cameras. It also means that the devices do not have to be placed near wall outlets, and means that power cables can be eliminated.

In PoE systems, a client device that receives power over an existing Ethernet network is called a powered device (PD). The device that delivers power to the PD is called power-sourcing equipment (PSE). Power consumption for the PD is limited to 12.95W, and PSE outputs are limited to 15.4W per RJ-45 port. Each PD can expect to draw a maximum continuous current of 350mA, assuming the

Ethernet link cable and physical layer device (PHY) transformers are well balanced.

To accommodate the voltage drop that takes place along a CAT-5 Ethernet link (up to 100m long), the IEEE standard specifies different power ratings for PD and PSE. The longer links exhibit a significant drop, which obliges the PSE to output more than the nominal 48VDC to maximize power available to the PD. Therefore, voltages as high as 57VDC can be seen anywhere along an Ethernet link.

Most PoE networks can be implemented with an endpoint or midspan PSE. An endpoint PSE integrates an Ethernet switch and power source in a single device, and is situated at the other end of the Ethernet link. This PSE type is the most convenient way to implement a PoE network, as power is already inline at the endpoint of an Ethernet link. Such Ethernet switches are sometimes referred to as having "inline power" (see **Figure 1**). Endpoint PSEs are ideal for new infrastructure deployments.

For existing Ethernet networks that cannot justify such a complete overhaul, power can be injected into the Ethernet link using the midspan PSE method. Midspan PSE provides power over the "spare pairs" in the CAT-5 cable—an approach that can be cost effective if only a few Ethernet devices need power. Such a case is the 4 to 24 ports in a local area that are part of a system comprising a larger multiport network (**Figure 2**).

Endpoint PSE differs from midspan PSE in its option to deliver power either combined with the signal over the same pair of wires, or over the spare pairs. In general, a PSE must be able to provide power over the signal pairs or the spare pairs, but not both.




Figure 1. For endpoint PSE and PD devices, power is delivered over the signal pairs.

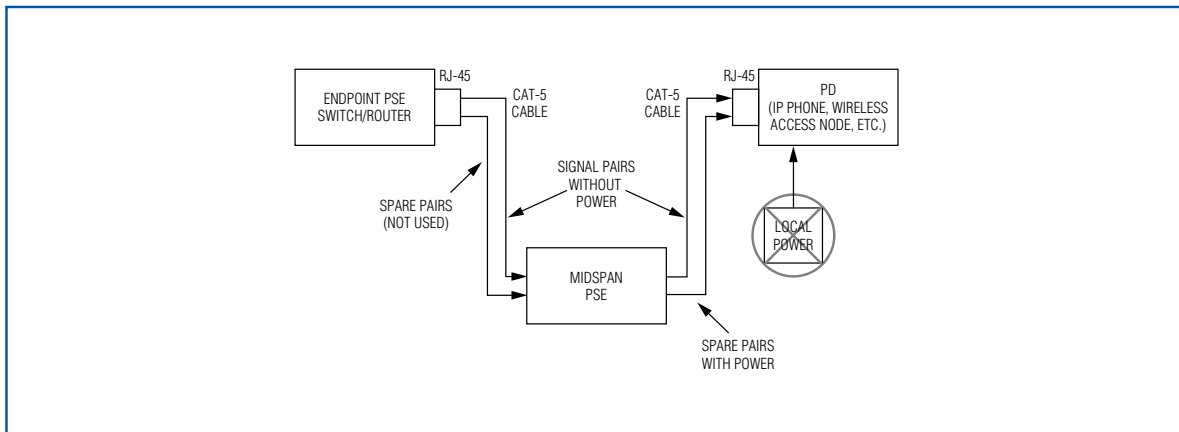


Figure 2. For midspan PSE and PD devices, power is delivered over the spare pairs.

Simple as they seem, such systems entail considerable design effort. They must include safeguards to ensure backward compatibility with devices that do not expect to see 48VDC on their Ethernet connection. The IEEE 802.3af standard covers backwards compatibility and, by including optional features for powering Ethernet networks, it also looks forward. This article covers what a designer should know in developing products destined to operate in new or existing systems—i.e., products expected to migrate towards Gigabit Ethernet or 1000BASE-T/TX.

What about Gigabit Ethernet?

Gigabit Ethernet works with endpoint PSEs, but not midspan PSEs, because it uses all four pairs within the CAT-5 cable for data transport. In contrast, 10BASE-T and 100BASE-TX use only two pairs for data (wires 1-2 and 3-6), leaving the spare pairs (wires 4-5 and 7-8) available for midspan power injection. To provide inline power for Gigabit Ethernet, therefore, endpoint PSE switches are required.

CAT-3 cable is supported by the IEEE 802.3af standard because it was originally used with 10BASE-T systems. To maximize signal integrity in new deployments, however, we recommend use of the highest rating of Ethernet cabling available (CAT-5e or CAT-6). This is because cabling infrastructures typically represent a ten-year investment. Gigabit Ethernet (1000BASE-T, specifically) requires CAT-5 cabling, but some applications using CAT-5 and Gigabit Ethernet switches have proven marginal. Consequently, the latest 1000BASE-TX standard requires CAT-6, while the original 1000BASE-T standard requires CAT-5.

Detection of PDs

When connected to its Ethernet links, the PSE must detect whether each of the Ethernet devices requires power. The PDs must therefore exhibit characteristics beyond those of a legacy Ethernet device. To accomplish this detection, the PSE makes V-I measurements while probing the signal wires with a current-limited voltage of 2.7V to 10.1V. **Table 1** lists the criteria a PD must have for

Table 1. For a valid-PD signature, all criteria below must be detected by a midspan or endpoint PSE.

Parameter	Conditions (V)	Minimum	Maximum
V-I slope (at any chord of 1V or greater)	2.7 to 10.1	23.75k Ω	26.25k Ω
Voltage offset	—	—	1.9V
Current offset	—	—	10 μ A
Input capacitance	2.7 to 10.1	0.05 μ F	0.12 μ F
Input inductance	2.7 to 10.1	—	100 μ H

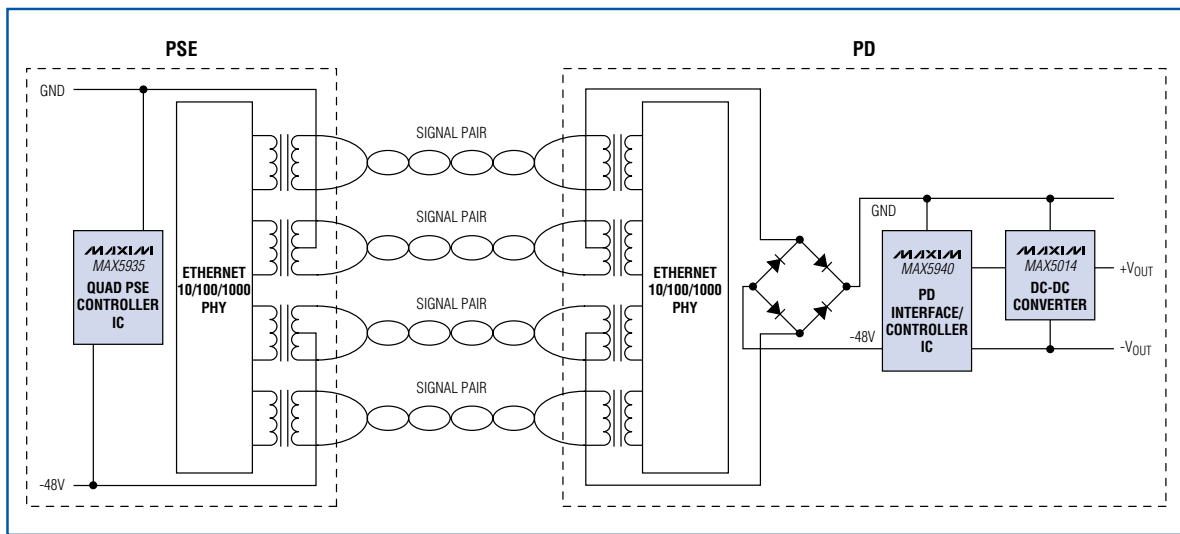


Figure 3. A PD operating with Gigabit Ethernet must be backward compatible with midspan PSE applications, and therefore receive power from an endpoint PSE switch.

detection as a valid PD. The 1.9V series offset allowed is a consequence of the diode bridges typically used to control voltage polarity. Two such bridges per PD are required, as the PD must be backward compatible with midspan PSE applications (Figure 3). The 10 μ A current offset is typically due to leakage within the PD. Table 2 lists another set of criteria, for which any detection fails an Ethernet device by making it an invalid PD.

Power classification of PDs

The driver that first started the movement to combine power with Ethernet networks was the voice-over IP (VoIP) telephone. Because so many other Ethernet devices are able to use this convenient source of power (RFID readers, chargers for PDAs, mobile phones, or even laptops), the IEEE 802.3af standard includes an optional feature called power classification. This allows the PSE to manage its power budget more closely. Table 3 lists the different power classes for which a PD can be provisioned, and their corresponding classification signatures.

To implement the optional power classification method, the PSE applies a probing voltage of 14.5V to 20.5V. In

response, the PD exhibits a signature (classification current), indicating back to the PSE the maximum power the PD can draw. That information enables the PSE switch to manage the maximum power it delivers to the connected PDs at any given time.

By selecting a proper PSE controller IC, you could implement another feature that is outside of the IEEE 802.3af standard: a hard limit on the PSE's output power per port. Unless the deployment administrator can guarantee that no PD will ever be swapped out for one that dissipates more power, the switch's expected power budget can occasionally be exceeded. In that case, the PSE will refuse to power the port unless the PD power classification is met.

Another feature handy in emergencies would be an ability for the PSA to prioritize which ports receive power first, or which ports should be disconnected first when the UPS or backup generator begins to run out of energy. The switch could then maintain power for the most important Ethernet ports. Such ports might include E911 telephones, badge readers for access, certain surveillance cameras or access points, or other revenue-generating data circuits.

Table 2. Detection of any criterion below by the midspan or endpoint PSE indicates that Ethernet device is an invalid PD.

Parameter	Conditions (V)	Range of Values
V-I slope	2.7 to 10.1	Either >45k Ω or <12k Ω
Input capacitance	2.7 to 10.1	>10 μ F

Table 3. Five classes for PD power classification and their classification signatures.

Class	Conditions (V)	Classification Current (mA)	PD Power Range (W)
0 (Default)	14.5 to 20.5	0 to 4	0.44 to 12.95
1	14.5 to 20.5	9 to 12	0.44 to 3.84
2	14.5 to 20.5	17 to 20	3.84 to 6.49
3	14.5 to 20.5	26 to 30	6.49 to 12.95
4 (Reserved for future use)	—	—	—

The presence of such fail-safe features within the PSE controller IC, either hardwired or software-configurable, can help manage the power budget during emergencies. Consequently, look for a software-configurable PSE controller IC.

Detecting disconnected PDs

After a PSE applies power to a PD, it must monitor the PD for a "maintain power" signature in accordance with the IEEE 802.3af standard. The PSE must also detect whether the PD has been disconnected. The standard defines both AC and DC methods for detecting a PD disconnect. Consider, for example, that a PD has been disconnected and a legacy Ethernet device immediately plugged into the same RJ-45 jack on the switch. If 48VDC power is not quickly disconnected after the PD is removed, the legacy device may be damaged.

AC-impedance measurements performed on a PD are generally more accurate than pure DC-resistance measurements. A small, common-mode AC voltage is sent down the Ethernet link simultaneously with the data signals and 48VDC. You then measure the AC current and calculate the resulting port impedance, which (if the PD has not been disconnected) should be less than 26.25k Ω . The frequency for this AC voltage must be between 1MHz and 100MHz. For the many other details pertaining to DC and AC methods of disconnect detection, the designer should consult the IEEE 802.3af standard. Regardless of the method chosen, the measurement must be made quickly and the power removed quickly thereafter.

Advanced features in silicon

Among the multiport PSE silicon chips now available, the most common are PSE controllers that control inline power to four ports. Look for I²CTM-compatible, serial-interface capability with programmable registers that provide the option of use with an MCU. Some of the advanced features residing in various operating modes are now important for emergency reasons, an importance that was magnified after 9/11/01.

Operational modes offered by the MAX5935, for instance, include automatic, semiautomatic, manual, shutdown, and debug mode. Automatic mode allows the device to operate without software supervision. Semiautomatic mode (on request) continuously detects and classifies a device connected to a port, but does not apply power to that port until directed by software. Manual mode, useful in system diagnostics, allows complete control of the device by software. Shutdown mode terminates all activity and turns off power to the ports. Finally, debug mode allows detailed system diagnostics by fine-stepping through the device state machine.

Figure 3 is a simplified example of a PoE system design, illustrating PSE and PD connections using Gigabit Ethernet. Because Gigabit Ethernet does not work with midspan power injection, the 100/10M Ethernet modes can only connect to an endpoint PSE switch. (The MAX5940 PD-interface controller does not need a diode bridge, yet it operates with one if required.) Today's PD-interface controller ICs (such as the MAX5941 and MAX5942) include a pulse-width modulation (PWM) controller, even though the PD usually includes a DC/DC converter.

Purchase of I²C components from Maxim Integrated Products, Inc., or one of its sublicensed Associated Companies, conveys a license under the Philips I²C Patent Rights to use these components in an I²C system, provided that the system conforms to the I²C Standard Specification defined by Philips.