

New Power for Ethernet—The LTC4255 Delivers

(Part 1 of a 3-Part Series) by Dave Dwelley

Introduction

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.¹ This promises a whole new class of Ethernet devices, including link-powered IP telephones, wireless access points, and PDA charging stations, which do not require additional AC wiring or external power transformers (“wall warts”). With about 13W of power available, small data devices can be powered by their Ethernet connection, free from AC wall outlets.

Modern Ethernet networks and traditional telephone systems share much in common. Both typically send data or voice over unshielded twisted-pair connections, and both are typically connected in a “star” configuration, where each terminal is connected to a central switch or hub. One significant difference, however, is that traditional phones are usually powered through the same wire as their “data” connection, whereas Ethernet devices require a local source of power. 802.3af changes this, by

allowing the central switch to provide 48VDC at up to 13W through the familiar RJ45 connector. Sophisticated detection and power monitoring techniques prevent damage to legacy data-only devices, while still supplying power to newer, Ethernet powered devices over the CAT-5 wire.

A device that supplies power is called a PSE (for Power Sourcing Equipment); a device that draws power from the wire is called a PD (for Powered Device). A PSE is typically an Ethernet switch, router, hub, or other network switching equipment that is commonly found in wiring closets or under desks where CAT-5 cables con-

verge. PDs can take many forms: digital IP telephones, wireless network access points, PDA or notebook computer docking stations, cell phone chargers, and HVAC thermostats are examples of devices that can draw their power from the network. Virtually any device that requires a data connection and can run from 13W or less can shed its AC power cord or batteries and operate off the RJ45 connector alone.

This article is the first in a three-part series on Powered Ethernet. This issue features Part 1, which covers the power details of the system, with a particular focus on the PSE and its characteristics. Part 2 will cover the PD in detail, while 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.

Delivering Power over Ethernet Cables

A CAT-5 Ethernet cable contains four unshielded twisted pairs of 24-gauge copper wire in a common sheath, with RJ45 connectors on each end. In a typical 10BASE-T or 100BASE-TX




Figure 1. Delivering power over existing Ethernet cables using the center tap of the transformer

Power Over Ethernet Glossary

- ❑ **PSE (Power Sourcing Equipment)**—usually a router or hub, but can also be a midspan
- ❑ **PD (Powered Device)**—any device that is powered over the Ethernet by a PSE: can be a phone, a WAP (Wireless Access Point) or even a PDA charger or an exit sign
- ❑ **Midspan**—a device that plugs in-line to convert a conventional router to a PSE; typically powers the spare pairs
- ❑ **Signal Pairs**—pairs 1-2 and 3-6 in CAT-5 cable
- ❑ **Signal Pairs**—pairs 4-5 and 7-8 in CAT-5 cable
- ❑ **PHY (Physical Layer Interface)**—the differential transceiver that transmits and receives data over the link

(10/100) network, two of the pairs (the “signal pairs”) are used for data transmission (one for transmit, one for receive) and two pairs (the “spare pairs”) are unused. 1000BASE-T (Gigabit over Copper) networks use all four pairs, and are compatible with most aspects of Powered Ethernet, although there are some incompatibilities and some aspects of the 802.3af standard do not explicitly support 1000BASE-T.

A PSE is required to provide a nominal 48V DC between either the signal pairs or the spare pairs (but not both)—see Figure 1. The power is applied as a common mode voltage difference between the two powered pairs, typically by powering the center-taps of the isolation transformers used to couple the differential data signals to the wire. Since Ethernet data lines are transformer-isolated at each end of the wire, this 48V potential difference between the transmit pair and the receive pair has no effect on the data transceivers on either end. The spare pairs can be tied together and powered directly (as shown in Figure 1), or they can be powered via transformer center taps in the same manner as the signal pairs if compatibility with 1000BASE-T is required.

The 48V supply used to power the line must be isolated from the PSE chassis ground to maintain the isolated link between the PSE and the PD. The IEEE defines two methods of isolation, named Environment A and Environment B. Environment A PSEs must isolate the 48V supply from the PSE chassis but need not isolate between adjacent ports, while the more stringent Environment B requires that ports be isolated both from the chassis and each other. In keeping with telecom conventions, the 48VDC supply is often referred to as a -48V supply; however, since the supply must be isolated from the chassis, which end is deemed to be “ground” is relatively arbitrary.

PSEs are physically located in one of two places: either integrated into data switch/router/hub devices, or as a standalone unit known as a

PSE Power Requirements

- ❑ Output is -44V to -57V (usually -48V), isolated from chassis—
Environment A: ports not isolated from each other
Environment B: isolated from chassis and port-to-port
- ❑ 15.4W (44V • 350mA) minimum power supply—current limit may drop as voltage rises.
- ❑ Turn on within 1s after PD is plugged in (single port only)
- ❑ Support $\leq 400\text{mA}$ loads for at least 50ms without current limiting—
Disconnect on overcurrent ($>350\text{mA}$) between 50ms and 75ms
Disconnect on undercurrent ($<5\text{mA}$) between 300ms and 400ms

“midspan” that connects in-line between an existing data switch and the PD. An integrated PSE/switch is allowed to drive either set of pairs, but typically will drive the signal pairs. A midspan is required to drive the spare pairs.

PSE Operation

A PSE is required to probe the cable for the characteristic PD signature before applying voltage to the wire. A valid PD signature consists of a 25k resistor with up to three diodes in series with it, and no more than 0.1 μF in parallel. The cable must be probed with voltages of less than 10V to minimize the chance of damaging a legacy data-only Ethernet device that may not be prepared to see 48V between its terminals. Only after a valid signature is detected may the PSE apply power to the wire.

After detecting a valid signature, a PSE may optionally check for a second PD classification signature that indicates the maximum power the PD will ever draw. This classification signature appears as one of several specific currents drawn by the PD when probed with a voltage between 15V and 20V. If the PSE opts to classify the PD, it can use the information to allocate power from a common power supply, or even deny power if it finds that the PD is requesting more than the PSE has available. The entire detection/classification/power up sequence must be complete within one second from the time the PD is first connected to the port.

Once the PSE has detected and optionally classified the PD and has

decided to turn on the power, it must provide between 44V and 57V (nominally 48V) to the appropriate pairs on the cable. The port must be able to supply at least 400mA for 50ms without current limiting, and must be able to supply 15.4W (44V • 350mA). As the port voltage rises, the PSE may reduce the current limit it allows, as long as the 15.4W power level is maintained. The 15.4W requirement allows for a PSE operating at the minimum voltage (44V) to supply the full 12.95W a PD is allowed to draw, plus the drop through a worst-case 20 Ω round-trip cable at the 350mA maximum continuous current. The port must limit output current to below 450mA at all times to protect against short circuits on the cable. If the PSE senses an overcurrent condition for more than 75ms, it must turn the power off.

Once the power is on, the PSE must keep it on as long as the PD presents a valid *power maintenance signature*. This power maintenance signature consists of two components, both of which the PD must exhibit: a minimum DC current draw of at least 10mA, and an AC impedance lower than 33k Ω at all frequencies from DC to 500Hz. The PSE can opt to monitor either or both components of this signature to determine if the PD is still present. If the PSE senses that the signature is invalid, it must wait between 300ms and 400ms before removing power from the line. The 300ms minimum prevents false disconnects caused by glitches on the line or sudden drops in the line voltage, and the 400ms maximum prevents a fleet-fingered technician

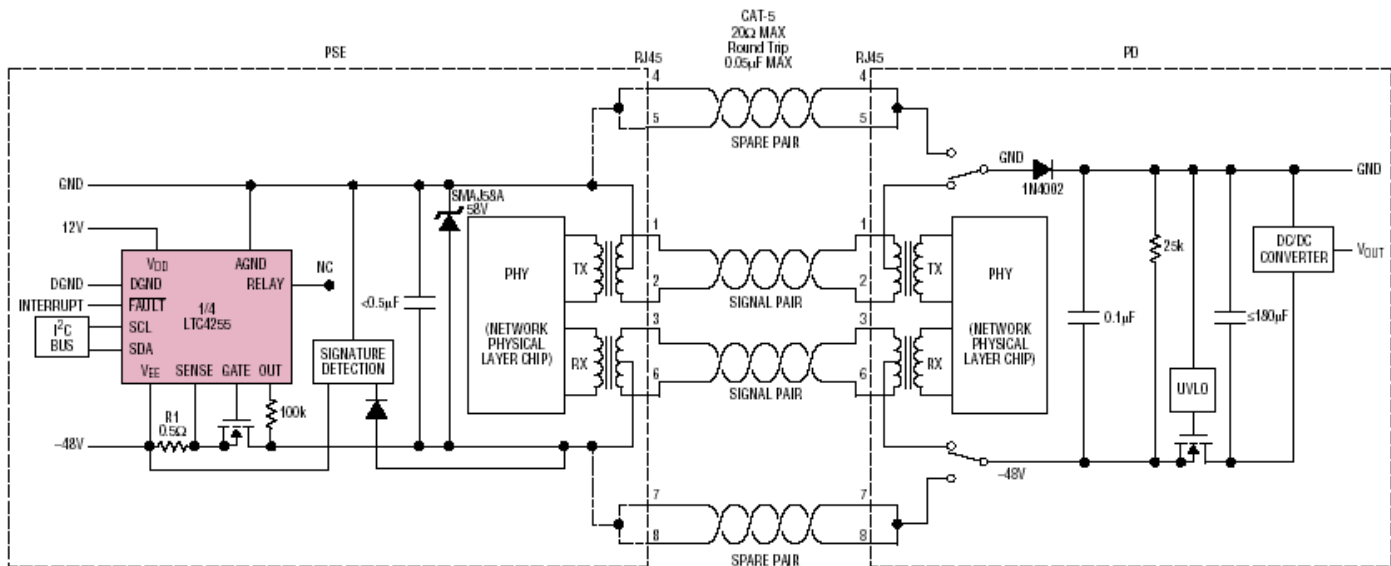


Figure 2. Power control circuitry using LTC4255 quad network power controller

from unplugging a valid PD and connecting a legacy device before the PSE has a chance to turn the power off.

LTC4255 Quad Network Power Controller

The LTC4255 is a quad -48V power controller, designed to implement the power path portion of a PSE device. It contains complete power management and switching circuitry for four channels, including -48V Hot Swap™ switching, current inrush control, current limit, and DC disconnect sensing for four ports. Internal status and control registers allow the LTC4255 to accept commands and report back status to the host system via the industry-standard two-wire I²C™ serial bus protocol. One LTC4255 channel, together with a standard differential data transceiver (commonly known as a “PHY”), a detection/classification circuit, and a couple of external components make a complete powered Ethernet port. The quad configuration of the LTC4255 makes it useful in multiport PSEs, such as powered Ethernet switches or hubs.

The primary function of the LTC4255 is to control the delivery of power to the PSE port. It does this by controlling the gate drive voltage to an external MOSFET (Figure 2) while monitoring the output current via

sense resistor R1 and the voltage at the OUT pin. This circuitry serves to couple the raw -48V supply to the port in a manner that meets the PD’s requirements while minimizing disturbances on the backplane.

When it receives an I²C bus command to turn on a port, the LTC4255 enters a timed startup mode where it powers up the PD in a mode that limits inrush current. The internal power control circuitry serves the gate drive to the external MOSFET to limit the port current, allowing the voltage at the port to rise in a controlled manner as the PD input capacitance charges. An internal timer controls the inrush duration, and foldback current limiting reduces the maximum current limit when the output voltage is below 30V , minimizing power dissipation in the external MOSFET. If the port reaches full voltage and the current draw drops below the current limit before the timer expires, the LTC4255 assumes the port turned on normally. It sets the Power OK bit in the status register and keeps the MOSFET gate turned fully ON until a disconnect or fault event occurs. If the port is still in current limit when the timer expires, the LTC4255 assumes there is something wrong with the PD, turns off the power and sets the corresponding fault bit in the status register.

Current Limit Protection

Once power has ramped up to its final value and the start-up timer has expired, the LTC4255 shifts to normal operation. In normal operation, the port current should never exceed the current limit level, I_{MAX} . The current limit circuit monitors the port current by watching the voltage across R1 and reduces the MOSFET gate voltage as needed to keep the current below I_{MAX} . When the current drops below I_{MAX} , the gate voltage is restored to the full value to keep the MOSFET resistance to a minimum.

If the port reenters current limit at any time after startup, a current limit timer starts. If this timer expires, the port is turned off and the fault bit is set in the corresponding power status register. The current limit timer is an integrating counter that decrements at a slower rate than it increments, preventing intermittent current limits from overheating the external power MOSFET.

DC Disconnect Detection

An additional current monitoring circuit trips when the port current drops below the minimum allowed level, signifying that the PD has been unplugged or has removed its power maintenance signature. If the current is still below the minimum when

continued on page 15


LTC4255, continued from page 11

the disconnect timer runs out, the port power is turned off and the corresponding status bit is set. The LTC4255 monitors the DC component of the Power Maintenance Signature only; additional circuitry is needed to monitor the AC component of the signature if required by the application.

Conclusion

The LTC4255 provides complete power control circuitry to switch 48V onto

Ethernet wires, greatly simplifying the design of the power path of PSE devices. An LTC4255, together with a standard quad PHY chip, a detection/classification circuit, and a handful of external components make four complete powered Ethernet ports. Fault protection, startup control, and disconnect sensing are all performed by the LTC4255, minimizing external circuitry. The I²C interface simplifies monitoring and control of the LTC4255 by a host system.

Part 2 of this series will cover the details of PD design and show how to put together the power receiving end of the link. 

Notes

¹ The 802.3af standard is still in draft form, and parts of the standard are still in flux. No product can yet claim full compliance, but compatible products are already available in advance of the final standard. For the latest information on the state of the 802.3af standard or on LTC products designed to meet the standard, contact the LTC Applications department.

For more information on parts featured in this issue, see
<http://www.linear.com/go/ltmag>

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

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.

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

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”).




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

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

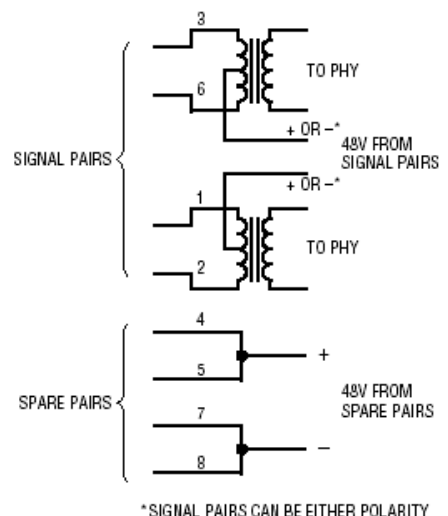


Figure 2. Deriving power from the cable

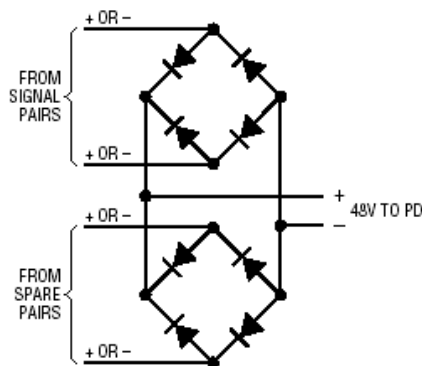


Figure 3a. Autopolarity input circuit




Figure 3b. Non-autopolarity input circuit

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

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

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 $\overline{\text{PWRGD}}$ 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. 

New Power for Ethernet—Detection and Classification (Part 3 of a 3-Part Series)

by Dave Dwelley

Introduction

The familiar RJ45 jack has been under the office desk now for better than a decade, and has served well as the data conduit to computers and devices with NICs (network interface cards). Now the IEEE 802.3af committee is extending the Ethernet standard to provide 48V DC power from these same RJ45 jacks, possibly relegating countless AC adaptors (wall warts) to the recycle bin.

This is a great leap forward, though it is not without pitfalls. For instance, what's to prevent legacy (non-ethernet-power-enabled) devices from burning up when a DC voltage that they were never designed to see appears on the wire?

The solution lies in the detection protocol that underlies 802.3af, which prevents a PSE (Power Sourcing Equipment) from applying power unless the device at the end of the wire is positively identified as an Ethernet-power-ready PD (Powered Device). This final article in a three-part series on Powered Ethernet covers detection and classification of PDs. Parts 1 and 2 covered PSEs and PDs, respectively, and are available at www.linear.com.

PD Detection

The fundamental defining characteristic of a PD is the 25k Ω signature impedance it must present to the wire when probed with voltages below 10V (Figure 1). Legacy network connections typically are open circuits or have a 150 Ω common-mode termination network between their terminals. 25k Ω is chosen to be well away from both of these impedances, to minimize the chance of false detection.

Before a PSE is allowed to apply power to the line, it must check for this signature resistance with a power-limited probing source. Typically, a PSE will use one of three probing strategies: force a small current on the line and




Figure 1. Simplified PD detection signature circuitry

measure the resulting voltage; force a small voltage on the line and measure the resulting current; or connect a voltage source to the line through a known resistance and calculate the result from the resistive divider.

“Linear Technology has several product families available that implement detection, classification, and power delivery for both PSE and PD systems.”

A PD need only present a valid signature when its terminal voltages are between 2.8V and 10V—a valid PD will pull a maximum of 550 μ A under these conditions. Under all circumstances, the PSE must limit the voltage while probing an open port to 30V or less, and limit the current while probing a short

to 5mA or less. To keep its internal circuitry from corrupting the signature, a PD typically includes a UVLO (Under Voltage Lock Out) circuit to disconnect everything but the signature circuitry from the line below 30V.

To be considered a valid signature, the PD must look like 25k Ω \pm 5% in parallel with 120nF or less of capacitance. The PSE, in turn, must accept a somewhat wider range of 19k Ω to 26.5k Ω to account for parasitic series and parallel resistances in the system. The PSE must reject anything below 15k Ω or above 33k Ω , or anything with >10 μ F across its terminals. Finally, if a PD decides for some reason it doesn't want power, it must make its signature resistance fall below 12k Ω or rise above 45k Ω to be sure it doesn't get detected. Figure 2 shows a diagram showing the resistance accept and reject bands.

PD Detection Details

The PD signature impedance is allowed to have a voltage offset of up to 1.9V (typically caused by up to two protection diodes in series), and a current offset of up to 10 μ A (typically caused by leakage in the PD). These terms complicate the PSE resistance measurement, since a single V-I point measurement will not account for these errors. As a result, the PSE is required to take at least two different V-I points, separated

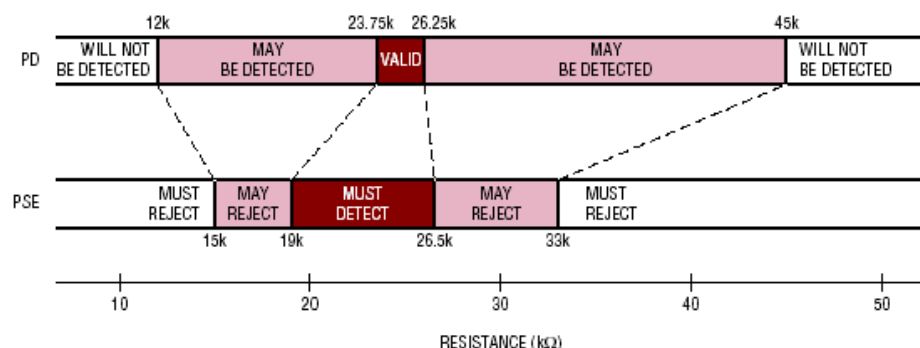


Figure 2. Detection resistance bands




Figure 3. Diode-OR on power inputs

by at least 1V at the PD. It then must calculate the difference between the two points to find the true resistive slope, subtracting out voltage and current offsets.

CAT-5 cable is also known as UTP, for Unshielded Twisted Pair. The data pairs are twisted together, providing effective protection from interference, but the pair-to-pair common mode signal that the PSE-PD link works on are not twisted or shielded, and provide virtually no rejection of interfering signals. Since CAT-5 cable is typically run in ceilings, walls, and other spaces where AC wiring is also present, 50/60Hz noise can be significant. The PSE must be able to successfully detect the 25kΩ signature with this line noise present, which it typically does by integrating several detection points or timing the detection points to be synchronized with the line frequency. The entire detection process must take less than one second from the time the PSE is first connected to the port to the time power is applied, limiting the total integration time during detection to a few hundred milliseconds.

We know from Part 2 of this series (*Linear Technology V12N4*, December 2002, pg. 9) that a PD has two possible power inputs: the signal pairs and the spare pairs. A PD must be prepared to accept power from either pair, which means it must present the detection signature simultaneously at each input. However, the PD is not allowed to receive power at both inputs simultaneously, even in unusual wiring configurations where live PSE ports appear at both inputs at the same time. The simplest way for a PD to meet these requirements is to diode-OR the two inputs into a single signature resistor (Figure 3). This has the advantage of

automatically selecting power from the active input, as well as disabling the signature at the second input when the first is powered.

The only time this scheme causes trouble is when both inputs are detecting simultaneously. In this case, the diodes alternately select one and then the other input as the detection probing signals rise and fall. This can corrupt the signature at each input to the point where neither PSE detects successfully. To avoid this problem, the 802.3af spec requires that a PSE powering the spare pairs remain idle for at least two seconds after a failed detection (open circuit excepted) to allow the signal pair PSE to retry successfully.

PD Classification

Once the PSE has successfully detected a PD, it can optionally ask the PD how much power it will draw. This second classification signature provides a way for the PSE to estimate how much power it needs to allocate to a particular port, since not all PDs draw the full 12.95W allowed by 802.3af. If the PSE has a large enough power supply to provide full power to all ports simultaneously, it can skip the classification step. If (more commonly) it has limited power available, it must keep track of how many PDs are connected and what their power classification levels are, and stop accepting PDs when the power budget is exhausted.

The classification signature is checked by forcing voltage across, or current into, the PD to push it into the 15V-to-20V classification signature band. Since we have already established that there is a valid PD

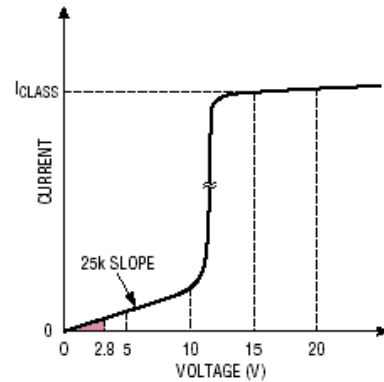


Figure 4. V-I characteristics of classification signature

connected, this additional voltage and current is unlikely to cause damage. In the classification signature band, the PD must behave like a constant current source, with a parallel impedance of 19kΩ or higher. The PSE measures this current and compares it to a set of fixed values that determine what class the PD falls into. Table 1 shows the available classes. The voltage source used during classification must be limited to 100mA to avoid damaging a malfunctioning PD, and it must not be connected for more than 75ms to keep PD power dissipation under control. Figure 4 shows how the PD current behaves as the system transitions from Detection mode to Classification mode.

Power Classes

PDs come in four classes, with a fifth reserved for future use. Classes 1 to 3 designate quarter, half, and full power PDs, respectively, as shown in Table 1. Class 0 is the default if a PD does not implement class signature circuitry; a 25kΩ signature resistor all by itself will land in the Class 0 current range.

continued on page 14

Table 1. IEEE 802.3af Power Classes

Class	PSE Minimum Output Power	PD Maximum Input Power	Nominal Classification Signature Current
0	15.4W	12.95W	<5mA
1	4.0W	3.84W	10.5mA
2	7.0W	6.49W	18.5mA
3	15.4W	12.95W	28mA
4	Reserved (Treat as Class 0)		40mA

The PSE interprets Class 0 as Class Unknown (which is usually treated as Class 3). Class 0 is typically reserved for extremely simple or low cost PDs where including classification circuitry is a burden. Class 4 is reserved for future use; a PSE that discovers a Class 4 PD should treat it as a Class 0.

The PD should place itself in the class that includes its peak current draw, even if the typical current draw is well below the class limits. For example, if a phone draws 2W when sitting idle but 8W when the ringer is ringing, it must class itself as a Class 3 (or Class 0) device even though typical power consumption is in the Class 1 range. A PD cannot change classes on the fly, since classification is only run once, right before power is applied.

Putting It All Together

Once a PSE has successfully detected and classified a PD, it then makes the decision whether to power it on. If the available power in the PSE is adequate to power the PD, the PSE turns on the power to the PD and begins monitoring the port for the Power Maintenance Signature, as described in Part 1 of this series (*Linear Technology* V12N3, December 2002, pg. 9). Once power is applied, the detection circuitry should be disabled, since the PD UVLO switch is now closed and the input impedance of the PD power supply circuitry typically swamps the 25k Ω signature impedance.

The PSE now has the whole picture: the detection sequence tells it that there is a real PD attached to the port; the classification routine tells it how much power that PD will draw so it can allocate its power supply resources

accordingly; and the Power Maintenance signature tells it that the PD is still present and operating normally. The PD, in turn, has a straightforward way to communicate to the PSE what it is, how much power it wants, and whether or not it wants that power to keep flowing. All of this goes on without affecting the data stream in any way. This allows both ends of the system to pick up or drop off power without affecting ongoing data transfers (assuming, of course, that the PD has an alternate power source). It also allows valid PSE and PD devices to be built with no data communications hardware at all—a significant cost savings for some applications. Most importantly, the detection protocol minimizes the chance of damaging a legacy Ethernet device, thus maintaining backward compatibility. 