

POWER OVER ETHERNET (PoE): TECHNOLOGY AND APPLICATIONS

Introduction

Power over Ethernet (PoE) is a natural offshoot of the need for networked peripherals to be powered independently of a dedicated source of power. Given that a network cable is already running to the peripheral, the question arises: Can we use that same cable to supply power to the device? The answer to that question is yes, we can use the data cable to supply power to devices, in addition to supplying data, with the right equipment. As a result, with Power over Ethernet we now have (speaking simplistically) a network with a structure composed of:

1. sources of data
2. sources of power (Power Sourcing Equipment, or "PSE"s)
3. wiring (usually Unshielded Twisted Pair Category 5 or better, or "Cat 5 UTP")
4. devices that can receive data and power (Powered Devices, or "PD"s)

The Power over Ethernet specification is not a stand-alone document; instead, it essentially bolts a whole new area onto the already-existing Ethernet specification. The areas of addition consist of a chapter that defines the structure, components, and interactions of the parts of a PoE system, as well as adding modifications to the Ethernet specification that interface with the operation of Power over Ethernet. PoE is an optional structure in Ethernet; data can pass in the system without PoE.

Basically, the Power over Ethernet specification consists of the following:

- definitions for how PSEs add power to Ethernet's existing 100 ohm balanced cabling system
- a description of a PD's load on the PSE and cabling
- a protocol that will detect devices requesting power from a PSE
- a method for classifying devices (PDs and others) according to their power needs
- a method for PDs and PSEs to dynamically negotiate and allocate power consumption
- a method for scaling back power supplied to a PD to a detect-only level when a PD no longer requests power

The Basic PoE Network

1) Source of data

The source of data in an Ethernet network is generally a switch or router, in the sense that these devices pass data between network devices. Switches and routers are not "sources of data" in the sense that they "create" data; that is the role of a computer, generally. And, in reality, data can flow either way.

2) Source of power

Power over Ethernet requires a device that puts power into the wiring of the network. Sources of power in a PoE network are usually either a PoE router or PoE switch (such devices are called "endspans"), or a power-injecting device (called a "midspan") between a standard router or switch and the power-receiving device at the endpoint. More on these power sources later; in PoE parlance they are "Power Sourcing Equipment", or "PSEs".

Figure 1: PoE Power Sources



3) Wiring

Wiring is the next component in our list, and Ethernet frequently runs in Cat 5 UTP (Category 5 Unshielded Twisted Pair) cabling, which has four pairs of wires, each pair twisted within itself. The twists minimize the effects of outside electrical interference; in addition, each pair is twisted with

a differing number of twists over a given length of cable to minimize crosstalk between pairs. But, as we will see, with only four wires needed for 10BASE-T and 100BASE-TX Ethernet signalling, four other wires are unused. For 1000BASE-T (Gigabit Ethernet) all wires are used for data.

4) Devices that receive data and power

These are any number of Ethernet devices, and in PoE-speak they are "Powered Devices", or "PDs". Devices receiving power over Ethernet can include:

- IP telephones
- IP cameras (where PoE can power pan, tilt, and zoom operation)
- wireless network access points
- RFID readers
- device servers, such as the LAVA Ether-Serial Links — PoE
- network routers or switches
- any other device attached to a PoE system, needing Ethernet data or not, whose power draw is within the parameters of the Power over Ethernet specification. Note that in this regard PoE can operate as a generalized power distribution methodology, with some of the added benefits of managed power.

Attempting to answer our initial question, "Can we use the network data cable to also supply power to network devices?" requires understanding the cable itself. As noted above, when used for 10BASE-T and 100BASE-TX Ethernet, network wiring has unused wires.

It didn't take long for people to start using those unused wires to carry power, but there was no standardized way of doing so until 2003, when the IEEE 802.3af-2003 Power over Ethernet standard was ratified. This standard defined the uses of Ethernet cabling for carrying power, and how that power was to be added, used, and managed. In fact, it also defined how the data-carrying wires in an Ethernet cable could also carry power. That standard was shortly followed with an extended PoE specification, IEEE 802.3at-2009. Both standards are described below.

IEEE 802.3af-2003

The initial standard, IEEE 802.3af-2003, can provide for Ethernet power of up to 15.4 W of DC power per device (with a minimum of 44 VDC at 350 mA). Depending on the type and length of cabling, some power losses can occur, down to an allowed 12.95 W. This

standard allows for PoE networks to run on Category 3 cabling and above.

IEEE 802.3af-2003 has a couple of operating modes ("Mode A" and "Mode B") that specify which wires are used for carrying power, and the modes in which endspans, midspans, and powered devices can operate. The table

"PoE Wiring" on page 8 delineates these characteristics.

IEEE 802.3af-2003 has three power class levels (plus a fourth unimplemented level) which are negotiated at initial connection (described in "Power Class Levels" on page 8).

IEEE 802.3at-2009

In 2009 the IEEE 802.3at-2009 standard came along. It subsumes the IEEE 802.3af standard, now retroactively calling it "IEEE 802.3at-2009 Type 1", and extends the capabilities of the standard, calling those extensions "IEEE 802.3at-2009 Type 2", "PoE Plus", or "PoE+". In practical terms, the newer standard upped the power available to provide up to 25.5 W. Some implementations conforming to this standard and using all four pairs of the Cat 5

cable can supply up to 51 W of power. A consequence of the added power requirements in this standard is that IEEE 802.3at networks run on Category 5 cabling and above; the Category 3 cabling allowed by IEEE 802.3af-2003 is no longer sufficient.

IEEE 802.3at-2009 has the same operating modes as IEEE 802.3af-2003, but has generalized their application more broadly

to endspans and midspans (described in "PoE Wiring" on page 8).

IEEE 802.3at-2009 four has power class levels (plus a fifth unimplemented level) which are negotiated at initial connection and can be dynamically modified thereafter (described in "Power Class Levels" on page 8).

Table 1: PoE Type 1 and Type2 Compared

	IEEE 802.3af—2003 (IEEE 802.3at—2009 Type1)	IEEE 802.3at—2009 Type2 (PoE Plus, PoE+)
Power at Powered Device (PD)	12.95W	25.50W
Maximum power supplied by the Power Sourcing Equipment (PSE)	15.40W	34.20W
Available power levels	Three discrete levels; one is set at power up	Four discrete levels: one is set at power up; ongoing 0.1W adjustments
Cabling	Category 3, Category 5	Category 5
Modes	Mode A – endspan Mode B – midspan	Mode A – endspan/midspan Mode B – endspan/midspan
Voltage range at PSE	44.0–57.0 V	50.0–57.0 V



	IEEE 802.3af—2003 (IEEE 802.3at—2009 Type1)	IEEE 802.3at—2009 Type2 (PoE Plus, PoE+)
Voltage range at PD	37.0–57.0 V	42.5–57.0 V
Maximum current	350 mA	600 mA per mode

PoE Operating Modes

The PoE standard can operate in a couple of modes (“Mode A” and “Mode B”), identified in the diagrams and table above. Earlier I made it sound like unused wires were needed for PoE; that’s not necessarily the case, although some earlier methodologies of transmitting power over Ethernet relied on the unused wiring pairs. PoE Type 1 and Type 2 can use either data pairs or non-data pairs, depending on Mode and Ethernet version.

In Mode A, the power is carried on the same wires as are carrying 10BASE-T and 100BASE-TX data: this mode is sometimes called “phantom power.” This mode is necessary for four-wire cabling and for PoE on 1000-Base-T (Gigabit Ethernet), which uses all eight wires

in Cat 5 cabling for data. Mode A implementations can have MDI or MDI-X polarities, as shown in “PSE Pinout Alternatives” on page 3.

In Mode B, the power is carried on the four unused (in 10BASE-T and 100BASE-TX) wires in a Cat 5 or similar cable. In a Gigabit Ethernet PoE implementation, phantom power used regardless of the wires used to carry power, as all wire pairs are used to carry data.

Powered Devices (PDs) must be able to use either mode. Power Sourcing Equipment (endspans or midspans) must only supply power when they detect a powered device is on the line.

The PSE determines whether a PoE connection will be a Mode A connection, a Mode B connection, or both Mode A and Mode B; the PD will indicate (with a specified resistance on its lines) whether it is PoE ready, and for what degree of power. The standard requires that a PD must be able of operating both in Mode A and Mode B.

Finally, the IEEE 802.3at-2009 standard, apart from offering more power, redefined the powered device’s operation to permit two powered devices to co-locate on one powered port, with one device using Mode A power and one using Mode B power.

Table 2: PSE Pinout Alternatives

Conductor	Mode A (MDI-X)	Mode A (MDI)	Mode B (all)
1	Negative	Positive	
2	Negative	Positive	
3	Positive	Negative	
4			Positive
5			Positive
6	Positive	Negative	
7			Negative
8			Negative

Figure 2, “10/100 Ethernet, Endspan, Mode A,” on page 4 below shows Power over Ethernet being delivered over a 10Base-T or 100Base-TX link, using the same lines as are being used for data (lines 1, 2, 3, and 6). Note that the other four lines (4, 5, 7, and 8), while existing in the cabling, are not used for delivering power or data to the PD in this case. The source of power in this diagram (the PSE) is internal in the switch/hub: this is an “endpoint” topology.

Figure 2: 10/100 Ethernet, Endspan, Mode A

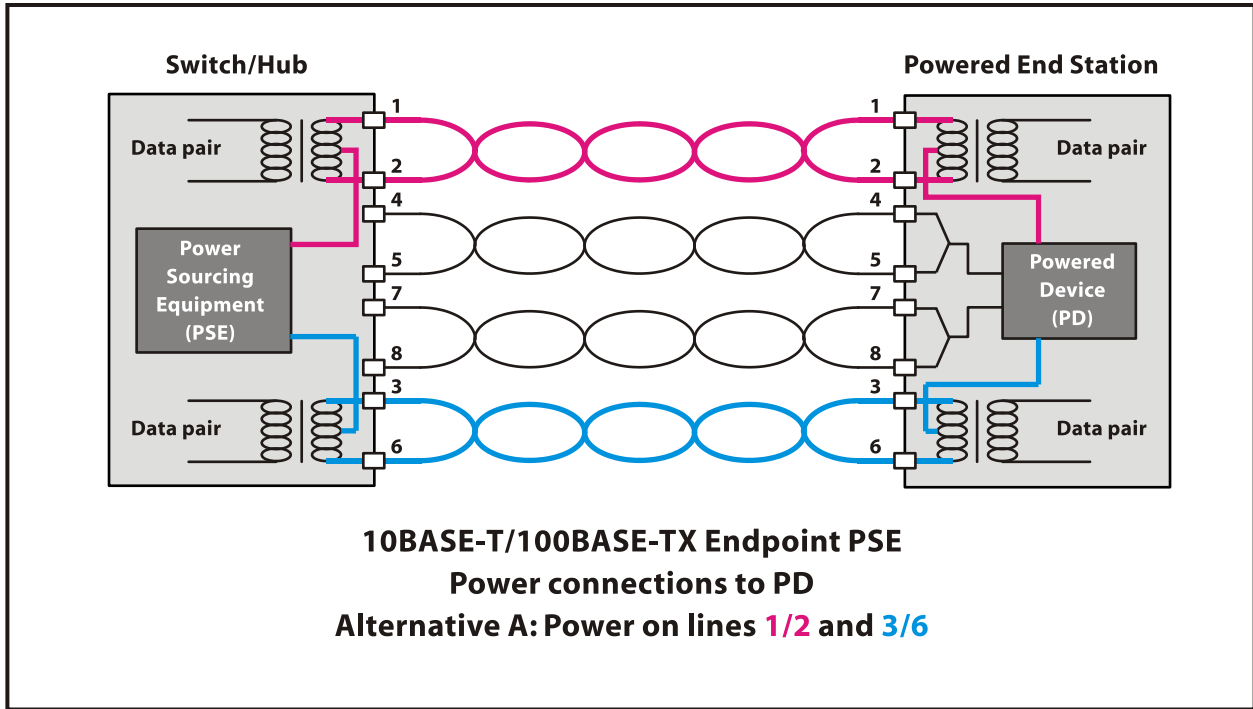


Figure 3, "10/100 Ethernet, Endspan, Mode B," on page 4 in this series, like the first, is a PSE endpoint delivering power over a 10Base-T or 100Base-TX link. The difference from the first diagram is that the power is being sent over lines unused for data transmission (lines 4, 5, 7, and 8). This diagram shows the PD as internal in the Powered End Station, but strictly speaking that is not necessary: the PSE could be powering a light bulb, with no data involvement whatsoever.

Figure 3: 10/100 Ethernet, Endspan, Mode B

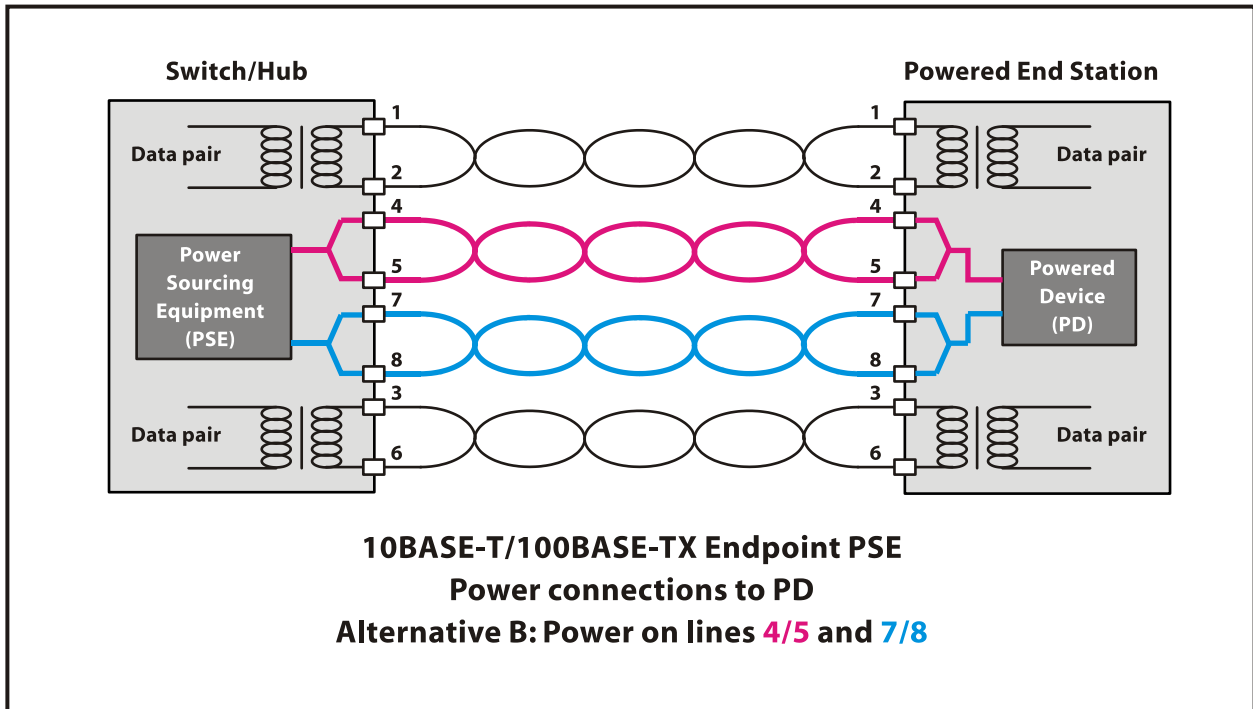


Figure 4, "1000 Ethernet, Endspan, Mode A," on page 5 shows an endpoint PSE sending power over a 1000Base-T (Gigabit Ethernet) network. Power is sent on lines 1, 2, 3, and 6 only, as with the 10Base-T or 100Base-TX networks in Figure 2, "10/100 Ethernet, Endspan, Mode A," on page 4, but as Gigabit Ethernet uses all eight lines for data transmission, lines 4, 5, 7, and 8 are also used, but just for data.

Figure 4: 1000 Ethernet, Endspan, Mode A

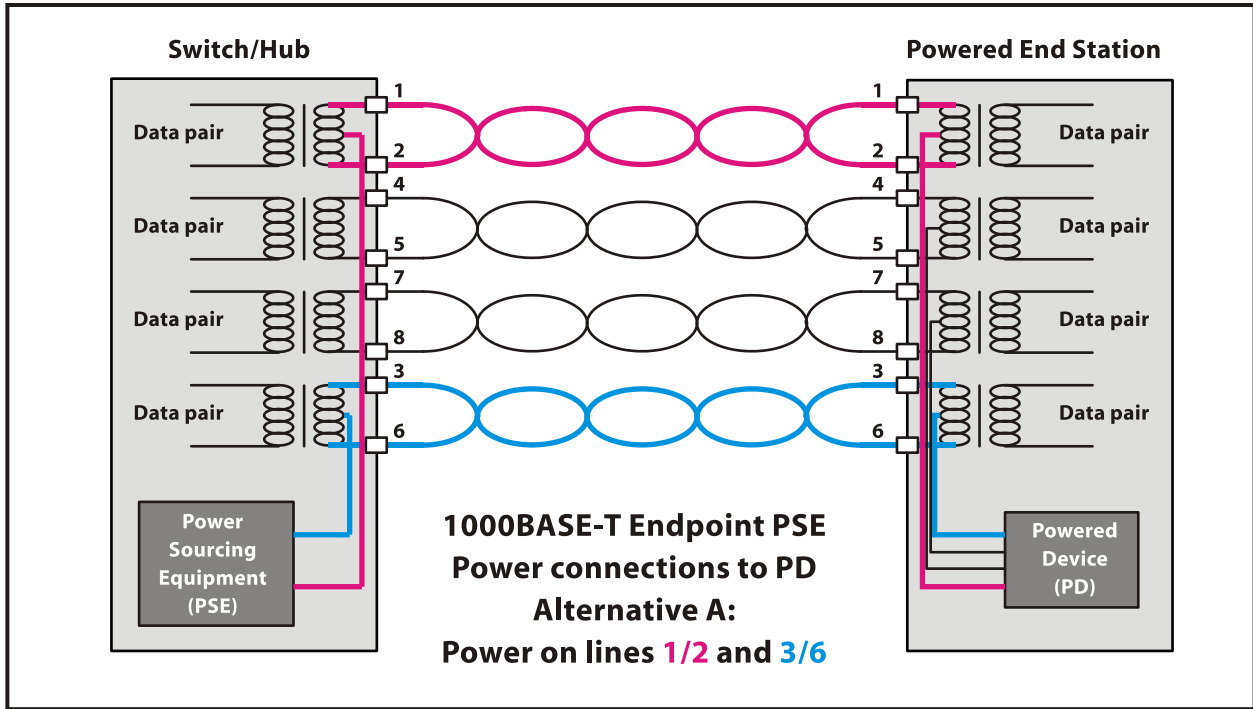


Figure 5, "1000 Ethernet, Endspan, Mode B," on page 5 is again a Gigabit Ethernet endpoint PSE, this time supplying power to the PD on lines 4, 5, 7, and 8. All lines are being used for data.

Figure 5: 1000 Ethernet, Endspan, Mode B

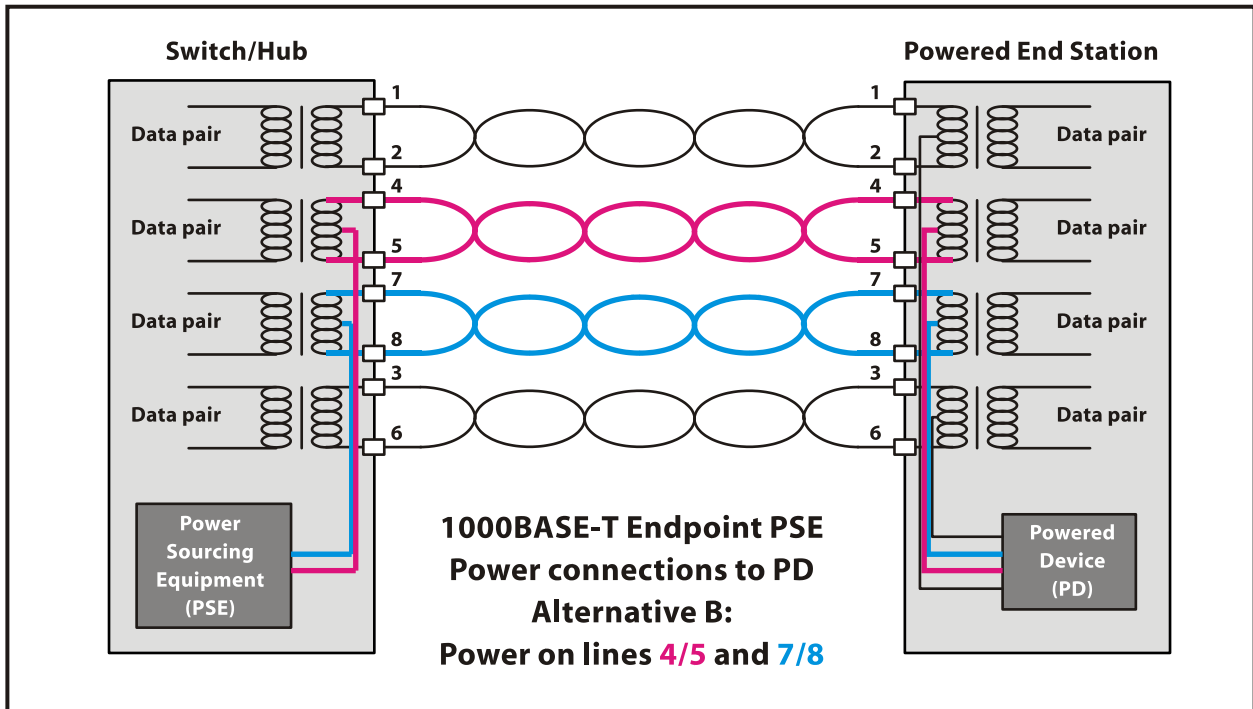


Figure 6, "10/100 Ethernet, Midspan, Mode A," on page 6 shows the PSE as a midspan device, inserting power into a 10Base-T or 100Base-TX segment. In this case, as the power is inserted on lines 1, 2, 3, and 6, the same lines being used for data, lines 4, 5, 7, and 8 are not being used for either power or data, so whether the midspan PSE even connects them is optional.

Figure 6: 10/100 Ethernet, Midspan, Mode A

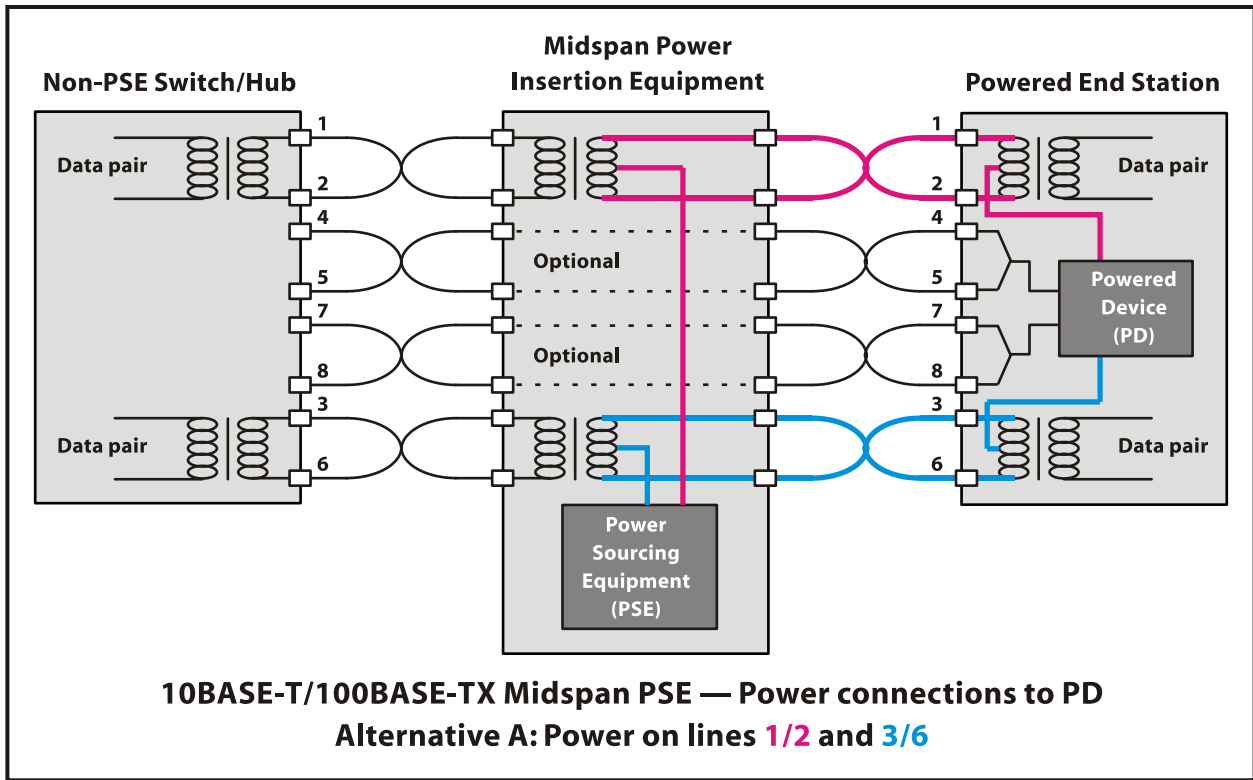
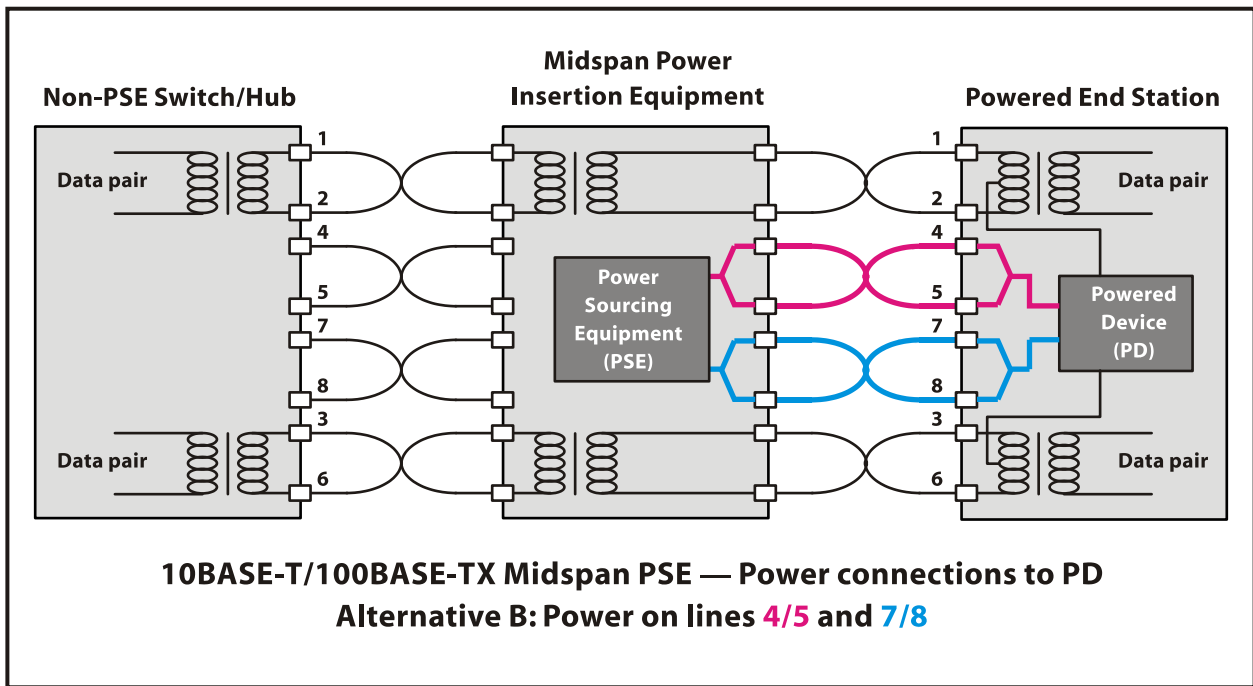


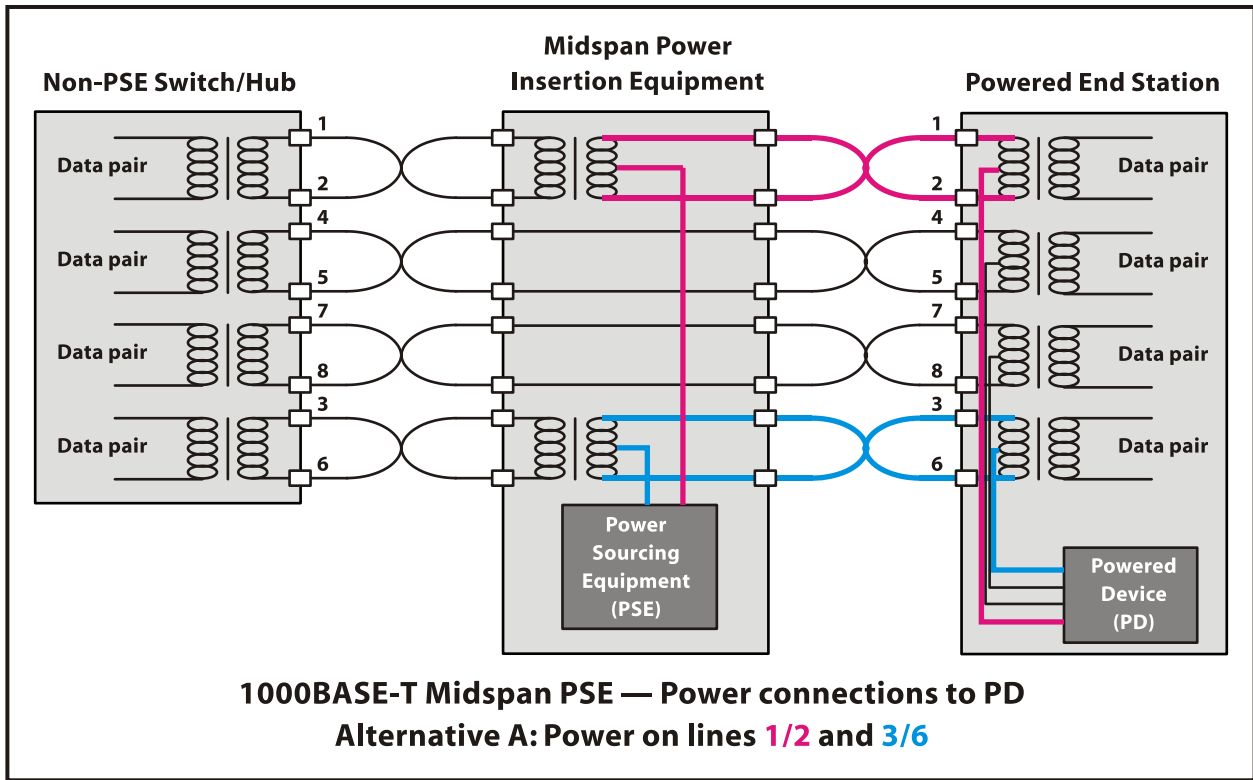
Figure 7, "10/100 Ethernet, Midspan, Mode B," on page 6 is a midspan PSE, inserting power on lines 4, 5, 7, and 8 (lines without data).

Figure 7: 10/100 Ethernet, Midspan, Mode B



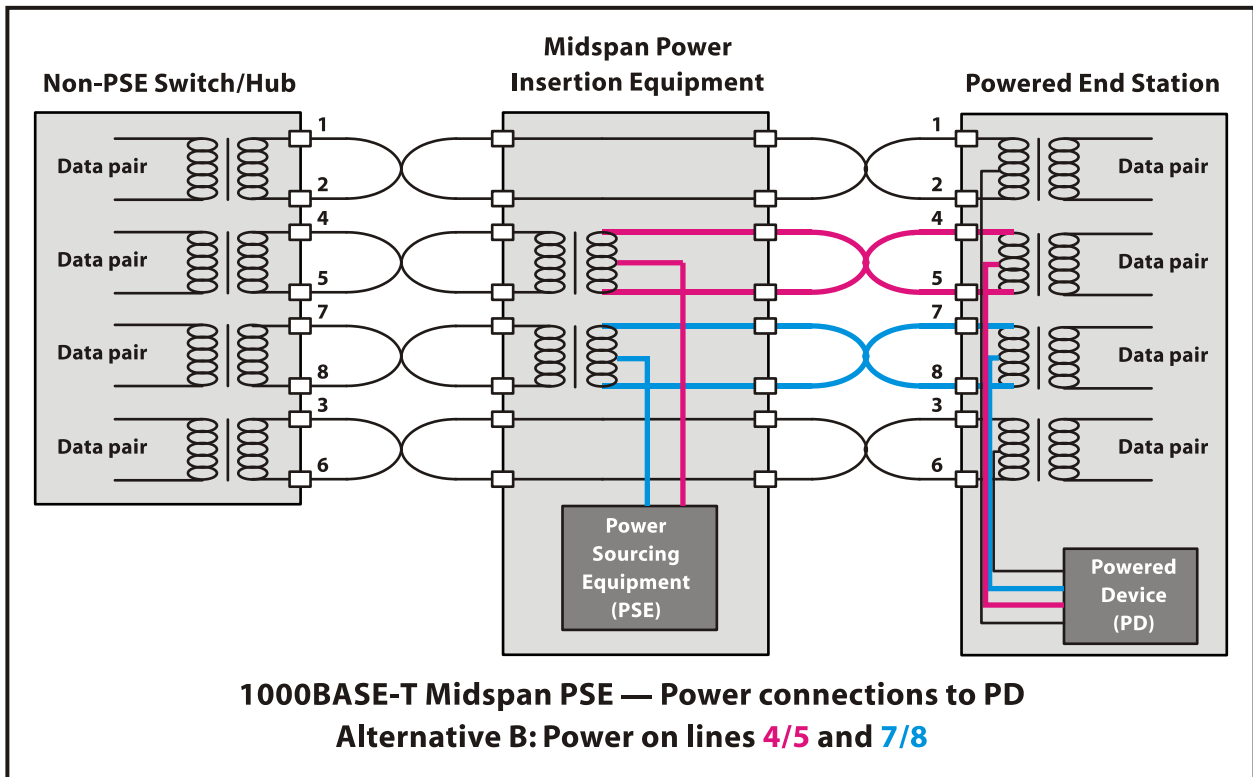
When a PSE midspan device inserts power on a Gigabit Ethernet connection, it must also pass all the lines through it, as Gigabit Ethernet needs all eight lines for data. Figure 8, "1000 Ethernet, Midspan, Mode A," on page 7 shows power inserted on lines 1, 2, 3, and 6.

Figure 8: 1000 Ethernet, Midspan, Mode A











Finally, midspan power insertion equipment can inject power on lines 4, 5, 7, and 8, while continuing to connect all eight lines to complete the required data connections, as shown in Figure 9, "1000 Ethernet, Midspan, Mode B," on page 7.

Figure 9: 1000 Ethernet, Midspan, Mode B



As can be seen, between Ethernet speeds, midspan versus endspans, and signal line modes, there are numerous possible combinations.

Figure 10: PoE Wiring

RJ-45 Pin #	Wire Color (T568A)	Wire Diagram (T568A)	10Base-T Signal 100Base-TX Signal	PoE
1	white/green		Transmit +	Mode A+
2	green		Transmit -	Mode A+
3	white/orange		Receive +	Mode A-
4	blue		unused	Mode B+
5	white/blue		unused	Mode B+
6	orange		Receive -	Mode A-
7	white/brown		unused	Mode B-
8	brown		unused	Mode B-

PoE Power Management

Table 3: Power Class Levels

Class	Usage	Classification current (mA)	Power range (Watts)	Class description
0	Default	0-4	0.44-12.94	Classification unimplemented
1	Optional	9-12	0.44-3.84	Very low power
2	Optional	17-20	3.84-6.49	Low power
3	Optional	26-30	6.49-12.95	Mid power
4	IEEE 802.3at Type 2 devices only	36-44	12.95-25.50	High power

How does PoE know which devices to power and how much power to give them?

Given that a network supplying power over Ethernet is likely to have devices that are not ready to accept power on their Ethernet connections, how does PoE prevent these devices from receiving potentially damaging power? The required power negotiation is handled in the PoE specification, by having the PSE look at a device to see whether it is presenting the appropriate PoE “detection signature”.

The detection circuit will tell a PSE whether there is a valid PD on the network segment, and ensure that the properly-timed response is provided. The detection circuitry also prevents two competing PSEs from powering the same network segment, and from incompatible polarities being used.

Detection leads to an initial determination of the validity of a PoE connection: whether a valid PSE-PD linkage is possible, and that there are no devices causing a PoE exculsion. Successful detection of a PSE-PD link then leads to classification, which happens in two

steps: physical layer classification and data link layer classification.

Physical Layer Classification means that the PSE classifies the PD’s power classification. This occurs before the PSE supplies power to the PD.

Once the Physical Layer Classification is sorted out, the PSE-PD connection powers up, with some common-sense thinking such as that the link powers up on the same line pairs as those that were used for the detection.

After the data link is established, Data Link Layer Classification is the means by which the PSE-PD link is managed, and this classification provides the ability for the PSE-PD connection to operate with dynamic power allocation.

Should the link become inactive, power will drop and a detect level of power will be maintained on the line.

The figures “Detection circuit: PSE detection source” on page 9 and “Detection circuit: Alternative PSE detection source” on page 9 show detection circuits that present the required detection signatures for PSE-PD connection. The diode D1 informs a PSE that another PSE has a reversed voltage, and presents a non-valid detection signature as a safety-preserving result.

The process of link classification is managed by the PSE, with appropriate fallback to lower power levels as the nature of the devices on the line segment are determined.

The detection signature is presented by the PD, as a resistance (R) calculated from two voltage (V)/current (I) measurements made during the interrogation by the PSE.

$$R_{detect} = \left\{ \frac{(V_2 - V_1)}{(I_2 - I_1)} \right\}_{\Omega}$$

Figure 11: Detection circuit: PSE detection source

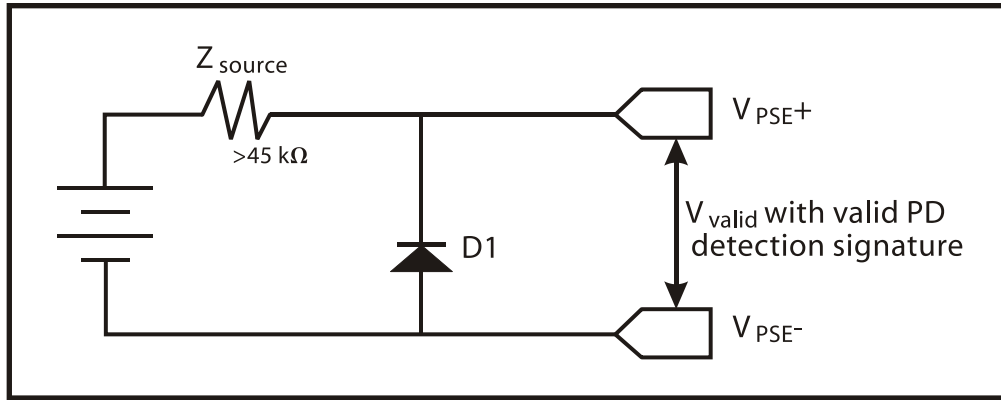



Figure 12: Detection circuit: Alternative PSE detection source



The permutations of Physical and Data Link layer classifications in the context of being either Type 1 or Type 2 PSEs or PDs generates the matrix "PSE and PD classification permutations" on page 9.

Table 4: PSE and PD classification permutations

Permutations				
PSE/PD Type	Physical Layer classification	Data Link Layer classification	PSE allowed?	PD allowed?
Type 2	2-Event	No	Yes	No
		Yes	Yes	Yes
	1-Event	No	No	No
		Yes	Yes	No
	None	No	No	No
		Yes	No	No
Type 1	2-Event	No	No	Yes
		Yes	No	Yes
	1-Event	No	Yes	Yes
		Yes	Yes	Yes
	None	No	Yes	No
		Yes	Yes	No

PDs are classified by current according to the following table.

Table 5: PD current classification

Measured I_{Class}	Classification
0 mA to 5.00 mA	Class 0
> 5.00 mA and < 8.00 mA	May be Class 0 or 1
8.00 mA to 13.0 mA	Class 1
> 13.0 mA and < 16.0 mA	May be Class 1 or 2
16.0 mA to 21.0 mA	Class 2
> 21.0 mA and < 25.0 mA	May be Class 2 or 3
25.0 mA to 31.0 mA	Class 3
> 31.0 mA and < 35.0 mA	May be Class 3 or 4
35.0 mA to 45.0 mA	Class 4
> 45.0 mA and < 51.0 mA	May be Class 4 or invalid class
NOTE: A Type 1 PSE may ignore I_{Class} and report Class 0.	

Advantages of PoE include:

- higher data rates than those available with AC powerline networking
- greater power supply than provided by USB (as well as longer cable runs)
- inexpensive cabling (existing Ethernet CAT5 UTP cabling can be used in most cases)
- standardized connectors (RJ-45)
- ability to use existing switches and routers, by injecting power using midspan devices
- independence from local power standards, assuming the endspan or midspan will operate in that power context
- centralized protection from surges and spikes when the endspan or midspan is protected

In addition, power over Ethernet has a number of powerful advantages over the small wall-plug power adapters often used with networked devices (often called “wall warts”). They include:

- no need for AC power at the location of the Powered Device
- centrally managed power and devices, where wall adapters are generally unmanaged
- surge and brownout protection that is not possible with “wall warts”
- the option for greater energy efficiency

This last point merits some detail. First, greater power efficiency comes from the fact that unlike traditional wall adapters, PoE Powered Devices do not need to be continually drawing power. Second, new and sophisticated PoE PSEs can also adapt to demand, power down devices as necessary, and adapt power output to suit varying cable lengths. Finally, when Power over Ethernet meets the newly developed IEEE 802.3az “Energy Efficient Ethernet” specification, additional power savings will be possible.

Owing to its simplicity, versatility, and use of existing cabling, PoE has the capability to play a large role in power distribution in the world of networked devices.

LAVA Ether-Serial Links: Power over Ethernet versions

LAVA makes four serial device servers that operate using Power over Ethernet. As a powering alternative, these units can also accept mains power with an optional power adapter.

