ethernet alliance

Overview of 802.3bt - Power over Ethernet standard

Lennart Yseboodt, Philips Lighting
David Abramson, Texas Instruments

April 2018

IEEE 802.3bt at a glance

The new Power over Ethernet 802.3bt standard is the third revision to the widely adopted IEEE standard that specifies low voltage power transfer to networked devices. The first IEEE PoE standard, 802.3af (2003), was able to provide 13 W to devices. This was increased to 25.5 W by 802.3at (2009). With 802.3bt, the amount of power available for devices increases nearly threefold, to 71.3 W, enabling a myriad of new applications. Publication of 802.3bt is expected in the second half of 2018.

The overview below shows the power levels defined by 802.3bt and how these relate to the existing PoE standard. All PoE devices (sources or loads) are interoperable, the only limitation being that a new high power PD (Powered Device) will not get the full power from an older or lower power PSE (Power Sourcing Equipment). The different power levels are defined as 'Classes'. There are four existing Classes, providing up to 25.5 W to PDs. This is being expanded with four new Classes, taking the power level higher, up to 51 W for Type 3 PDs and up to 71.3 W for Type 4 PDs.

	Type 3 (802.3bt)						Type 4 (802.3bt)	
	Type 1 (802.3af)			Type 2 (802.3at)				
PSE	Class 1 4 W	Class 2 7 W	Class 3 15.4 W	Class 4 30 W	Class 5 45 W	Class 6 60 W	Class 7 75 W	Class 8 90 W
	2-pair only (Type 1 & 2) 2-pair or 4-pair power (Type 3 & 4)				always 4-pair power			
PD	Class 1 3.84 W	Class 2 6.49 W	Class 3 13 W	Class 4 25.5 W	Class 5 40 W	Class 6 51 W	Class 7 62 W	Class 8 71.3 W

This higher power level is made possible by providing current through all four twisted pairs in a network cable. The current standard (IEEE 802.3-2015/2018) only permits power transfer over two out of four pairs. For assigned Class 5 and higher, power delivery requires 4 pairs to be used. The new Type 3 / Type 4 PSEs also support existing PDs, and may also use 4 pairs to deliver power to these PDs, resulting in cable losses being halved.

In addition to more efficient power delivery and far more available power, 802.3bt has a number of new features:

Short MPS (Maintain Power Signature) allows PDs to achieve a much lower standby power compared to the existing standard. The minimum standby power has been reduced to 1/10[th] of what the current standard allowed (20 mW versus 200 mW). This enables IoT applications to be powered with PoE and have acceptable standby performance.

Autoclass is an optional classification technique that allows the PSE to account for the resistive losses in the cable and optimize the power allocation based on a reference power measurement. This allows the PSE to power more ports from a limited power supply budget.

Power demotion allows a PSE that cannot meet the power demand of a PD to provide it with a lower Class. The PD can then operate in a limited mode using the available power.

Mandatory classification is a change to the PSE rules for hardware based classification. Where Type 2 PSEs were not required to support full hardware classification, and could in stead use

LLDP (a protocol over the Ethernet data link) to provide full power to PDs, it is now mandatory for Type 3 and Type 4 PSEs to fully support hardware classification, which leads to a more robust system. LLDP is still used by PDs to fine tune their power demand.

LLDP extensions are a set of new fields in the LLDP protocol definition, that allow information exchange about: 4-pair capability of PDs, Autoclass, the maximum amount of power a PSE has, timed power down of a PD, measurements of voltage/current/power/energy, and the exchange of electricity price information.

About this document

This document provides a generic overview of how the 802.3bt IEEE Power over Ethernet standard works, with emphasis on the new elements compared to the existing standard (802.3-2015). The information herein is based on P802.3bt Draft 3.3.

While every effort is made to provide correct information, there is always the inherent risk to unintentionally deviate from what the specification actually says. This whitepaper is intended to help understand the structure and some of the more challenging parts of IEEE 802.3bt, and should not be used in lieu of the IEEE 802.3bt standard.

> Each major section in the document highlights the relevant portions in the IEEE 802.3bt standard by listing these in a green box at the beginning of the section.

The focus of this white paper is to explain operation of single-signature PDs, as this is the most common PD implementation. Dual-signature operation is substantially different on most levels and will be covered in another Ethernet Alliance white paper. Unless specifically mentioned otherwise, references to "PSE" and "PD" without a specified Type, refers to Type 3 and Type 4 devices.

This document is prepared by Lennart Yseboodt (Philips Lighting, Chief Editor for IEEE P802.3bt DTE Power via MDI over 4-Pair Task Force) and David Abramson (Texas Instruments, Comment Editor for IEEE P802.3bt DTE Power via MDI over 4-Pair Task Force). The views presented in this document should be considered the personal views of the authors and not a formal position of the IEEE.

Contents

List of Figures

1. Introduction

Power over Ethernet is a technology to transfer a limited amount of low voltage DC power over a standard communication cable, concurrent with the Ethernet data flow. While a number of technologies exist that offer power transfer under the name "PoE", by far the most widely used is the IEEE 802.3 Power over Ethernet standard.

While superimposing a DC voltage on a network cable isn't hard — all it takes is a voltage source and a couple of data transformers with a center tap after all — doing so in a safe and inter-operable way is not trivial. Great effort was spent to ensure that all devices compliant to the IEEE Power over Ethernet standard work with one another. The first PoE standard published in 2003 and counted 58 pages. New standards in 802.3 typically live in their own chapter, called a 'Clause' in the 802.3 document. The 802.3af amendment created Clause 33, titled "Data Termi-nal Equipment (DTE) Power via Media Dependent Interface (MDI)." This PoE standard supported up to 13 W for a PD, and introduced the majority of the concepts that exist in the standard today.

The need for higher power soon kicked off the successor project, 802.3at, published in 2009, counting 92 pages. The "DTE Power enhancements" Task force completely replaced Clause 33 with new text that describes both the "af" standard and the new "at" standard. This amendment (and Clause 33) were titled "Data Terminal Equipment (DTE) Power via the Media Dependent Interface (MDI) Enhancements." While consistent with the original, this title did not have a clear link to the popular term "Power over Ethernet" used in the industry. We'll come back to that later on.

802.3at introduced the concept of a device "Type". Type is a container of a number of key specifications for a given device, and is part of the device's identity. Type 1 refers to the 802.3af specification, whereas 802.3at introduced Type 2, which increased the power level to 25.5 W for the PD. Type 1 and Type 2 are fully interoperable, a Type 2 PSE will happily power a Type 1 PD, and a Type 1 PSE will power a Type 2 PD. Obviously, there is no way to get more than 13 W out of a Type 1 PSE, but the system is designed such that a Type 2 PD will recognize it is being powered by an (older) Type 1 PSE, and can only draw 13 W.

In 2013, once again fueled by a need for more power, the 802.3bt study group began the task of defining power delivery over 4 pairs. Currently at Draft 3.3 in IEEE-SA Sponsor Ballot, the standard numbers 175 pages. The additional complexity is in part caused by the need to be interoperable with the existing standard, which in itself is already an expanded version of the original. Instead of trying to describe the new Types and 4-pair power in Clause 33, we chose to create a new Clause for the new standard. The new Clause still provides full interoperability between all the device Types. The two Clauses were also given a more sensible title, Clause 33 is now titled "Power over Ethernet over 2 Pairs" and Clause 145 is titled "Power over Ethernet".

With 802.3 now claiming the term "Power over Ethernet" clearly in the title of the specifying Clause, a first step is taken to tackle the industry confusion around the term. The Ethernet Al-liance has developed a certification program for PoE products based on IEEE 802.3-2015. Further information about this program may be found at `https://ethernetalliance.org/poecert/`. Work is currently underway to create the second generation of the certification program to ad-dress products based on the IEEE 802.3bt standard. Organizations interested in this effort should contact poe_certification@ethernetalliance.org.

802.3bt introduces two new Types: Type 3, which supports up to 51 W delivered to a PD, and Type 4, which supports up to 71.3 W delivered to a PD. Consider Type 3 the replacement/upgrade of

the existing standard, while doubling the amount of power for the PD. Type 4 is the high power Type, able to provide more than 51 W to the PD. Type 4 PSEs have a higher minimum output voltage, which reduces the required current to provide these high power levels. Type 4 PSEs are also required to regulate their output power not to exceed 99.9 W for longer than 4 seconds to stay within the definition of a Limited Power Source (LPS). Type 3 PSEs may not exceed LPS power for any time duration.

The Type 3 specification's first four Classes are identical to those in Clause 33 (Type 1 and Type 2) to ensure complete interoperability between the "old" Type 1 and Type 2 devices and the "new" Type 3 and Type 4 devices. As such, 802.3bt aims to eventually obsolete the older Power over Ethernet standard in Clause 33, while retaining full backwards compatibility.

2.   System overview

A Power over Ethernet system consists of a PSE (Power Sourcing Equipment) connected to a PD (Powered Device) through a link section. "Link section" is the 802.3 term for the cable connecting the PSE to the PD. The PSE delivers power to a PD, by applying a voltage to the network cable. The PD consumes the power sourced by the PSE.

When a PD is connected to a PSE, the general order of events is as described in Table 1. IEEE 802.3 Power over Ethernet defines mandatory checks that prevent voltage being applied unless a valid PD is connected.

power flow →

| PSE
Power Sourcing Equipment | link section | PD
Powered Device |

aka. the network cable

Figure 1: Power over Ethernet major components

There are two kinds of PSEs. "Endpoint PSEs" are Ethernet data switches that have built in PSE functionality. A "Midspan PSE" on the other hand is inserted between a switch on one end, and to the PD on the other and injects power onto the link segment. This allows PoE to be added on any Ethernet link that does not have power by inserting a Midspan PSE in between.

Table 1: Typical startup procedure for a Power over Ethernet system

	Connection between the PSE and the PD is made.
Detection	The PSE will probe the pairset (or pairsets) for a PD. Power is applied only when a valid PD detection signature is found. PSEs that are configured to apply 4-pair power will perform detection on both pairsets.
Connection check	If the PSE is configured to apply 4-pair power, it will perform a connection check to discover if the connected PD is a single-signature PD or a dual-signature PD. Connection check does not necessarily happen after detection, it can happen before, concurrently with, or after detection. Connection check is exclusive to Type 3 and Type 4 PSEs that are 4-pair capable.
Classification	After the PSE has performed successful detection (and connection check if applicable) it proceeds with classification. Classification performs two functions: the PSE discovers how much power the PD wants, and it informs the PD how much power it may draw.
Inrush	The first phase of delivering power to the PD is called inrush. During inrush the PSE actively limits the amount of current that flows. The purpose of this is to prevent excessive inrush current to flow when the PD starts up.
Powering	This is the operational phase, the PSE monitors the current/power draw and disconnects the PD in case it exceeds a set limit. The PSE also checks if the PD becomes disconnected, in which case power is removed from the port.
Power removal	Power may be removed for a number of reasons, for example the PD is disconnected from the PSE, the PSE no longer has power available, the PD draws excess power, or there is a short circuit in the cable.

2.1. Power Interface

An important concept in the standard is the Power Interface, or PI. The PI is the physical point where the cabling connects to the PSE or to the PD. In essence it is where the two modular connectors touch. All specifications in the standard apply at the PI, anything inside the PSE or PD is considered "implementation specific".

The PI consists of 8 individual contacts, that correspond to the 8 conductors in a communication cable. Two conductors form a pair (a twisted pair in the cable), both always at the same nominal voltage. There are four pairs. A pairset consists of a positive and a negative pair. There are two pairsets, which are called Alternative A and Alternative B for the PSE, and Mode A and Mode B for the PD. The relations between these concepts is shown in Figure 2.
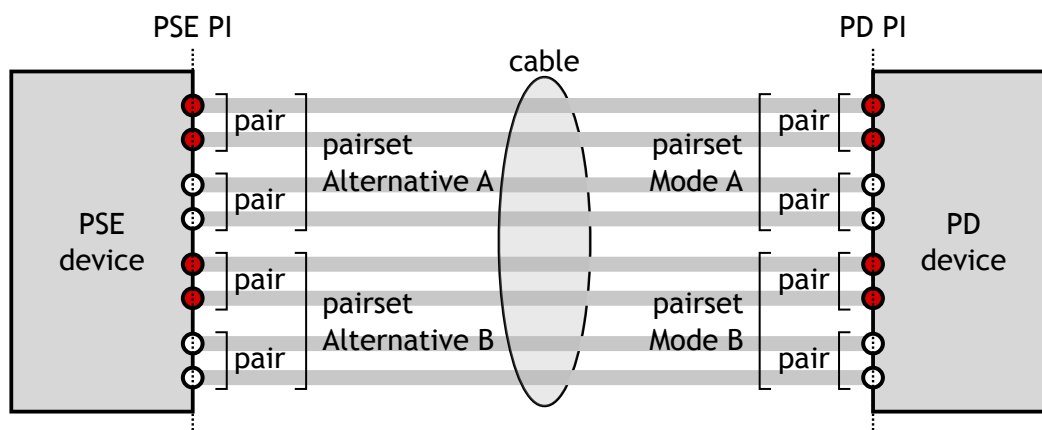
Figure 2: Power Interface (PI)

2.2. Class

The power Class is a key concept in the Power over Ethernet standard. The Class defines the maximum power being sourced or drawn in the system. There are eight Classes, numbered 1 through 8, where 8 represents the highest power level. Class 0 (an alias for Class 3) is defined in Clause 33, but deprecated in Clause 145. Many of the device parameters and requirements depend on the (assigned) Class. For more details see Section 5.

2.3. Device Types

The device Type determines the major static characteristics of the PSE and the PD. An overview of the key characteristics of each Type is shown in Table 2 for PSEs and Table 3 for PDs.

Type 1 and Type 2 PSEs only provide power over 2 pairs. A Type 3 PSE can either be capable of only 2-pair power, in which case it can support up to Class 4 (25.5 W PDs), or it can be capable of supplying 4-pair power in which case it can support up to Class 6 PDs (51 W). A Type 4 PSE is always capable of 4-pair power and can support up to Class 8 PDs (71.3 W). Note that PSEs are not required to be able to support the maximum Class of the given Type, the actual rules are:

PSE Type	Class support requirements
Type 1 PSEs	must be able to support at least Class 1, and may support up to Class 3.
Type 2 PSEs	must be able to support Class 4.
Type 3 PSEs	must be able to support at least Class 1, and may support up to Class 6.
Type 4 PSEs	must be able to support at least Class 7, and may support Class 8.

In addition to the highest Class a PSE supports, it also supports all the lower Classes (eg. A Type 2 PSE supports Class 0, 1 , 2, 3, and 4).

Unless explicitly mentioned otherwise, this document refers to Type 3 and Type 4 devices. In the standard this is similarly handled: in Clause 33, "PSEs" or "PDs" refers to Type 1 and Type 2 devices, where in Clause 145 it refers to Type 3 and Type 4 devices, unless explicitly mentioned.

Table 2: PSE Type overview

		802.3 Clause	Max power	Max Class	Hardware classification	Autoclass	Short MPS
Type 1	2-pair only	33	15.4 W	3	optional	✖	✖
Type 2	2-pair only	33	30 W	4	optional	✖	✖
Type 3	2-pair only	145	30 W	4	✔	optional	✔
Type 3	4-pair capable	145	60 W	6	✔	optional	✔
Type 4	4-pair capable	145	90 W	8	✔	optional	✔

Table 3: PD Type overview

	802.3 Clause	Max power	Classes	Autoclass
Type 1	33	13 W	0, 1, 2, 3	✖
Type 2	33	25.5 W	4	✖
Type 3	145	51 W	1 to 6	optional
Type 4	145	71.3 W	7, 8	optional

2.4. Cable / link section requirements

The Power over Ethernet standard defines several important terms related to the DC resistance of the link section. Correct operation is guaranteed for a "maximum pairset DC loop resistance" of up to $12.5\,\Omega$. But what does that mean? Let's explain some terms used in the standard.

conductor A network cable consists of 8 individual conductors, with a typical wire gauge between 26 AWG and 22 AWG.

twisted pair A twisted pair consists of two conductors twisted per a certain twist ratio. Power over Ethernet uses the common mode of these pairs to transfer DC power, without disturbing the data. The two conductors on a given pair are at the same nominal voltage and conduct current in the same direction.

pairset A pairset consists of two twisted pairs that together allow the transfer of power. One pair in the pairset is at the positive voltage, the other pair is at the negative voltage. When we say "supply power over 2-pair", this means that one pairset (thus, 2 pairs) is being used. The pairsets have specific names:

PSE name	PD name	cable pair colors
Alternative A	Mode A	green and orange
Alternative B	Mode B	blue and brown

R_{Ch} This parameter is defined as the "maximum pairset DC loop resistance" with a fixed value of $12.5\,\Omega$. A loop resistance is the cable resistance the electrical current would encounter from the source to the load and back through a single pairset. Note that this parameter represents the worst-case value.

R_{Chan} This parameter is the "actual DC resistance from the PSE PI to the PD PI and back". It represents the actual resistance in the system. In case of 2-pair powering, it is the loop resistance of the active pairset, with a maximum value of $R_{Ch} = 12.5\,\Omega$. In case of 4-pair powering it is the loop resistance of both pairsets in parallel, with a maximum value of $R_{Ch}/2 = 6.25\,\Omega$. This parameter is mostly used for specifications dealing with single-signature PDs.

- For 2-pair mode, as shown in Figure 3, $R_{Chan} = R_{Pair1} + R_{Pair2}$.

- For 4-pair mode, as shown in Figure 4, $R_{Chan} = (R_{Pair1} \,||\, R_{Pair3}) + (R_{Pair2} \,||\, R_{Pair4})$.

$R_{Chan\text{-}2P}$ This parameter is the "actual DC pairset resistance from the PSE PI to the PD PI and back". It represents the actual resistance in the system of a particular pairset, and is mostly used for specifications dealing with dual-signature PDs. $R_{Chan\text{-}2P}$ is thus either $(R_{Pair1} + R_{Pair2})$ or $(R_{Pair3} + R_{Pair4})$

Refer to Section 4 for information on single-signature PDs and dual-signature PDs.

2.5. 2-pair and 4-pair powering

Two powering modes exist: 2-pair mode, as depicted in Figure 3, and 4-pair mode, as depicted in Figure 4. Type 1 and Type 2 PSEs exclusively provide power in 2-pair mode. Type 3 PSEs can be constructed to be capable of 4-pair powering, or restricted to 2-pair powering. Type 4 PSEs are capable of providing 4-pair power.

The *capability* to provide power over 4-pair is distinct from what the PSE will actually do. In order to be permitted to deliver 4-pair power, a number of preconditions must be met (see Section 4.3). Furthermore, for Class 1 through Class 4 power, a PSE is permitted to either use 2-pair or 4-pair powering. It may also freely switch between 2-pair and 4-pair mode when it is providing power. A PSE that assigned Class 5 or higher power is required to provide power over 4 pairs at all times. See Section 8 and Table 14.

When power is supplied over only 2 pairs, the power can be supplied over Alternative A (pair 1 and pair 2), or over Alternative B (pair 3 and pair 4).


Figure 3: Supplying power over 2 pairs

Figure 4: Supplying power over 4 pairs

3. Detection

> PSE: 145.2.6 PSE detection of PDs
> PD: 145.3.4 PD valid and non-valid detection signatures

Detection is the most important function in a PoE system. It determines if the remote equipment connected to a PSE is capable of receiving power (in other words, it determines if it is a PD). The detection method itself has not changed from that used by Type 1 and Type 2 devices, but some of the requirements related to detection have been expanded.

Let's start with an overview of the how detection works. During detection, the PSE measures the effective resistance of the PD using at least two voltage/current points. Note that effective resistance is not the same as taking a single (absolute) resistance measurement. By taking the difference of the two measurements the resistive slope can be determined, which is used for detection evaluation. This differential resistance measurement allows static voltage offsets in the PD to be accounted for.

$$R_{detect} = \frac{V_{detect1} - V_{detect2}}{I_{detect1} - I_{detect2}}$$

An effective resistance of 25 kΩ signifies that the remote device is a PD and is capable of receiving power. Figure 5 shows the effective resistance requirements for the PSE and the PD.

The PSE may source a voltage and measure current, or it may source a current and measure voltage. The two measurement points must be at least 1 V apart and occur in the range of 2.8 V to 10 V at the PSE PI. As detection is used to check if the remote device is capable of accepting power, it must be done in a manner that is safe to all network equipment. Thus, the PSE cannot produce a voltage higher than 30 V and cannot source more than 5 mA of current during detection.

For Type 1 and Type 2 devices, the detection requirements are specified "at the PI" of the device, as only one pairset would ever be detected/powered at a time, since Type 1 and Type 2 PSEs provide power over 2 pairs only. For Type 3 and Type 4 devices, both pairsets can be powered and thus each pairset must be independently detected before applying power on that pairset. Under no condition can a PSE power a pairset, unless that pairset has a valid detection signature. Detection for Type 3 and Type 4 devices is specified to apply to pairsets.

In order to have a "valid detection signature" a pairset must have:

- A resistive slope between 23.7 kΩ and 26.3 kΩ, which is shown as the green triangle meets the R_{detect} requirement, shown in Figure 6. The PD must meet this slope over the entire PD detection voltage range.

- A voltage offset (V_{offset}) of no more than 1.9 V. Because the detection resistor (and the PD circuitry) is behind a rectifier, some voltage is lost before it reaches this resistor. For this reason, the PSE accepts the resistive slope to be shifted right by up to 2 V on the PSE side (V_{os}).

- The PD is not allowed any current offset, however the PSE must still accept a current offset (I_{os}) of up to 12 µA. Any current offset would cause the slope line in Figure 6 to shift upward.

- An input capacitance between 50 nF and 120 nF.

- An input series inductance less than 100 µH.

The voltage offset can be calculated from the resistance measurement as follows:

$$V_{offset} = \frac{V_{detect2} \cdot I_{detect1} - V_{detect1} \cdot I_{detect2}}{I_{detect1} - I_{detect2}}$$
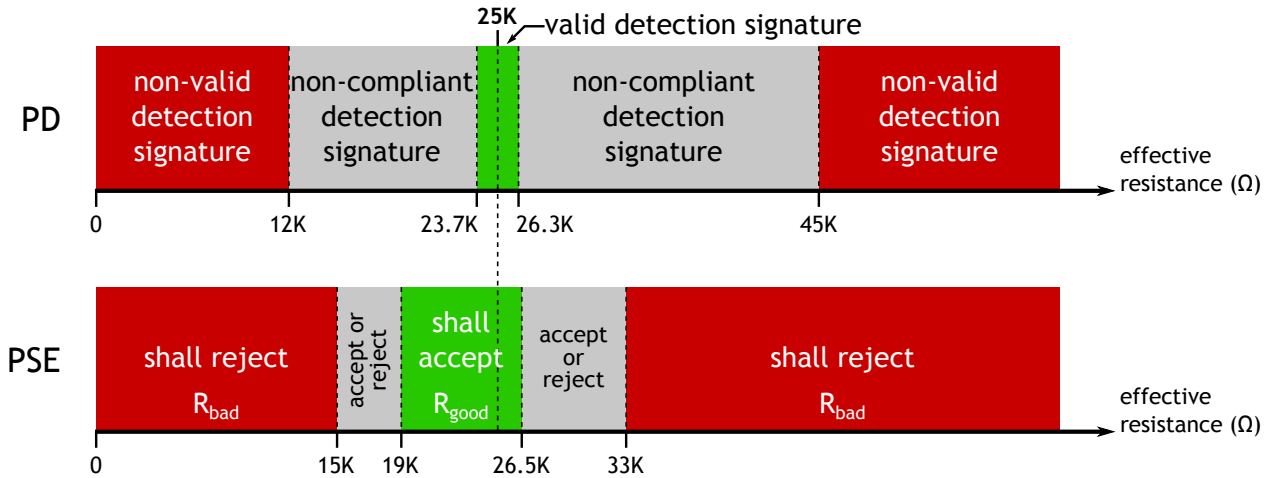
Figure 5: PSE and PD effective resistance detection parameters
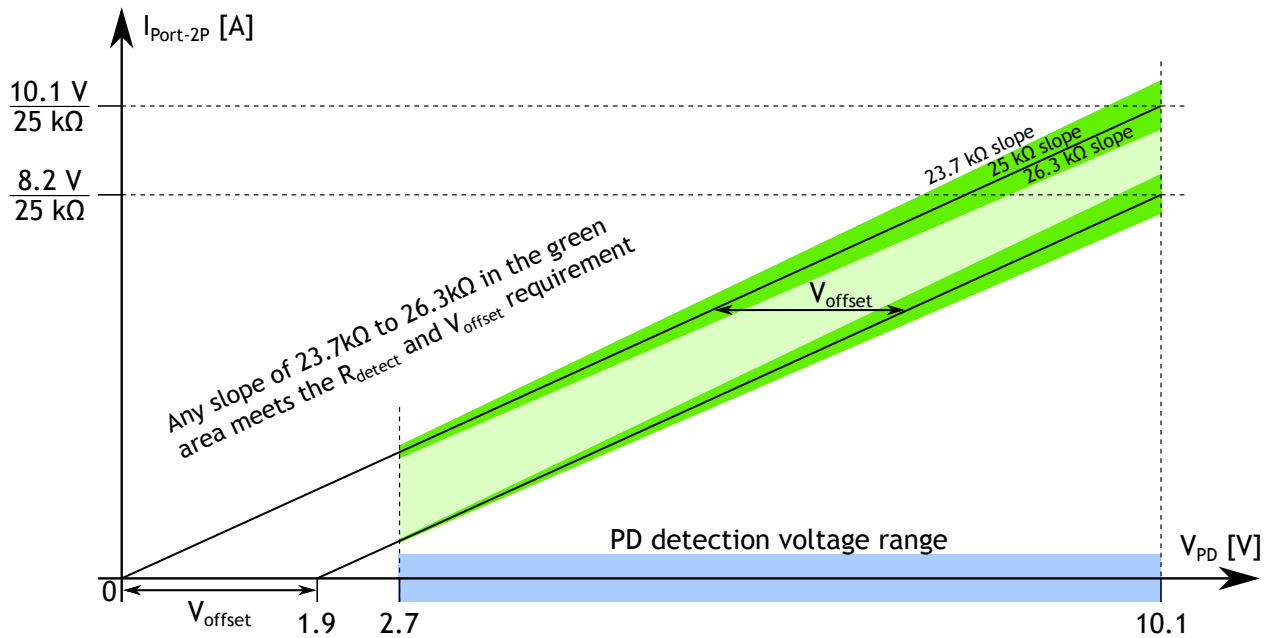
Figure 6: Voltage vs current plot of a valid PD detection signature

4. Connection check and PD signature configuration

> PSE: 145.2.6.1 Connection check requirements
> PD: 145.3.5 PD signature configurations

Two basic kinds of PDs exist: single-signature PDs and dual-signature PDs. The requirements for these two kinds of PDs are quite different, both for the PD itself, and for the PSE that is powering these devices. This document focuses on single-signature PDs and PSEs powering single-signature PDs.

Connection check is the mechanism a 4-pair capable PSE uses to probe the PD to find out if it is a single-signature configuration, a dual-signature configuration, or an invalid PD. The standard does not specify a specific method to determine the PD signature, this is considered implementation specific. An implementation can make use of the defined properties that single- and dual-signature PDs have to make its determination. These core properties are explained in Sections 4.1 and 4.2.

Detection and connection check exist for different purposes and should not be confused. Detection serves to identify an attached device as being a PD that wants to receive power. Connection check determines if the PD is single-signature or dual-signature. The result of connection check also helps to determine if a PD may be powered over 4 pairs (see Section 4.3) and how this PD will be treated by the PSE.

4.1. Single-signature PD configuration

A conceptual diagram of a single-signature PD is shown in Figure 7. The essence of a single-signature PD is that both Modes are connected to the same supply rail (through the rectifiers). The detection and classification mechanism as seen through either Mode A or through Mode B is the same one. A single-signature PD has one or more electrical loads, all sharing the same supply rail.

Figure 7: Concept diagram of a single-signature PD

The requirement defining a "single-signature PD", is tricky and subtle. The standard specifies all requirements at the PI, and for this particular requirement, an indirect property had to be found that indicates both pairsets are connected together.

The requirement is twofold:

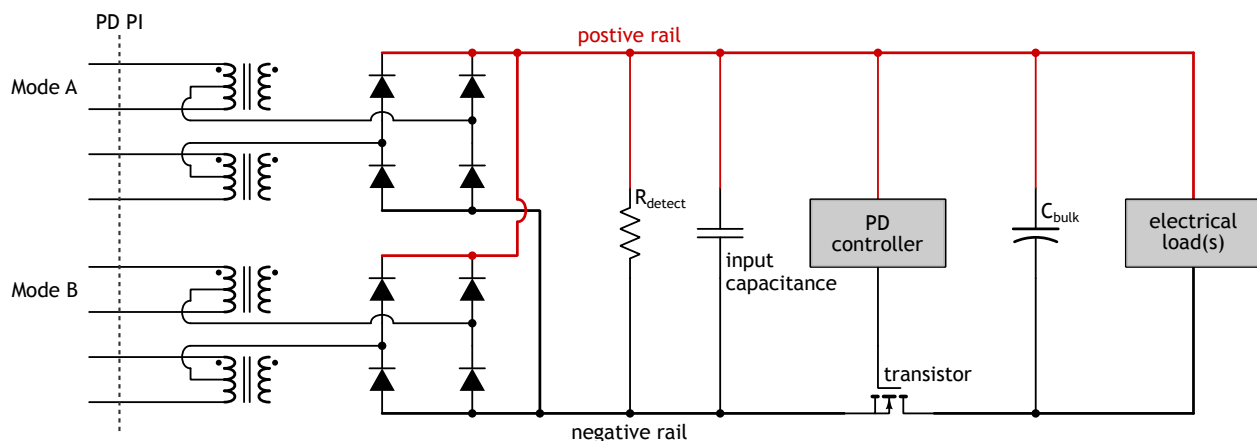- "A single-signature PD shall present a valid detection signature on a given Mode when no voltage or current is applied to the other Mode" …

- … "and shall **not present a valid** detection signature on the given Mode when any voltage in the range of 3.7 V to 57 V is applied to the other Mode or any current greater than 124 µA is applied to the other Mode."

The first requirement is simple: have a valid detection on either pairset, when the other pairset is left alone. A valid detection signature, plotted as current versus voltage, is shown in Figure 6. A PD with a resistive slope between 23.7 kΩ and 26.3 kΩ, which is shown as the green area in Figure 5, meets the R_{detect} requirement.

The second part of the requirement specifies what must happen to the detection signature when either a voltage is applied to the other Mode, or a current is injected. A PD that has connected both positive outputs from the rectifier into a positive rail and both negative outputs from the rectifier into a negative rail will no longer show a valid detection signature when either a voltage of at least 3.7 V to 57 V is applied, or a current of at least 124 µA is injected. The manner in which the detection signature loses validity is completely different though. This requirement is symmetric: it applies to Mode A when Mode B is manipulated with voltage or current, and applies also to Mode B when Mode A is manipulated with voltage or current.
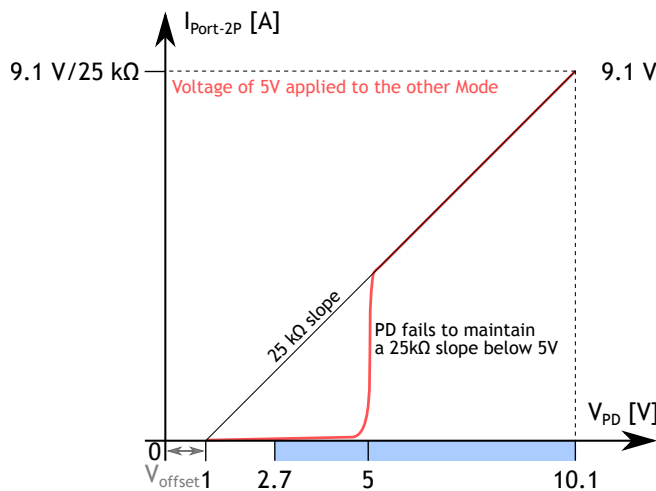
Figure 8: Voltage vs current plot of a PD with 5 V applied to the other Mode
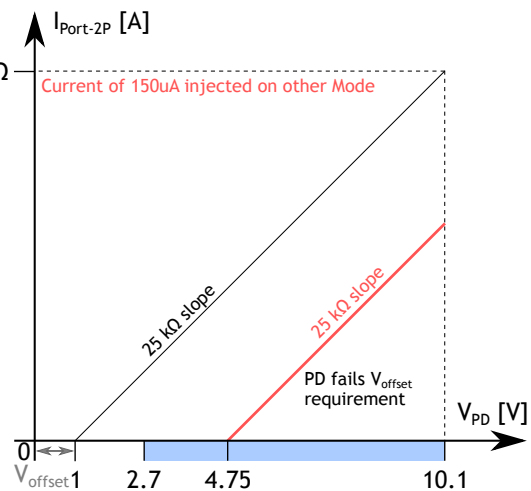
Figure 9: Voltage vs current plot of a PD with 150 µA applied to the other Mode

Take as an example a PD that has a $V_{offset} = 1$ V and a 25 kΩ detection resistor. Without any influence on Mode B, the Mode A voltage/current slope would look like the black slope lines in Figure 8 and Figure 9. However when a voltage of 5 V is applied to Mode B, the resulting voltage/current plot will become similar to the red plot in Figure 8. The voltage source on Mode B prevents any current from flowing until the voltage on Mode A exceeds that of Mode B. As soon as PD PI voltage on Mode A exceeds the PD PI voltage on Mode B the current follows the expected

25 kΩ slope. The PD is said "not to have a valid detection signature" because it fails to maintain the correct R_{detect} slope below the voltage that is applied to Mode B.

In Figure 9 a current of 150 µA is injected into Mode B. The resulting voltage vs current trace is shown as the red trace, which appears as a parallel line offset below the 25 kΩ slope. The injected current causes a voltage offset to appear over the detection resistance. Only when the voltage of Mode A exceeds that of Mode B, current can be begin to flow. In this case, while the slope is correct, the PD does not have a valid detection signature because it exceeds the V_{offset} requirement. This is an example, note that other methods can be implemented to make this determination.

4.2. Dual-signature PD configuration

A conceptual diagram of a dual-signature PD is shown in Figure 10. The essence of a dual-signature PD is that both Modes connect to an individual PD controller. The detection and classification mechanism as seen through Mode A or through Mode B is a different one. This means that even if Mode A is powered, it is still possible to do detection and classification on Mode B. This would not be possible with a single-signature PD. Figure 10 shows a dual-signature PD with completely isolated loads. Other configurations are also possible, but out of scope for this overview.

The requirement that defines a "dual-signature PD" is much easier to understand: "A dual-signature PD shall present a valid detection signature on a given Mode, regardless of any voltage between 0 V to 57 V applied to the other Mode".
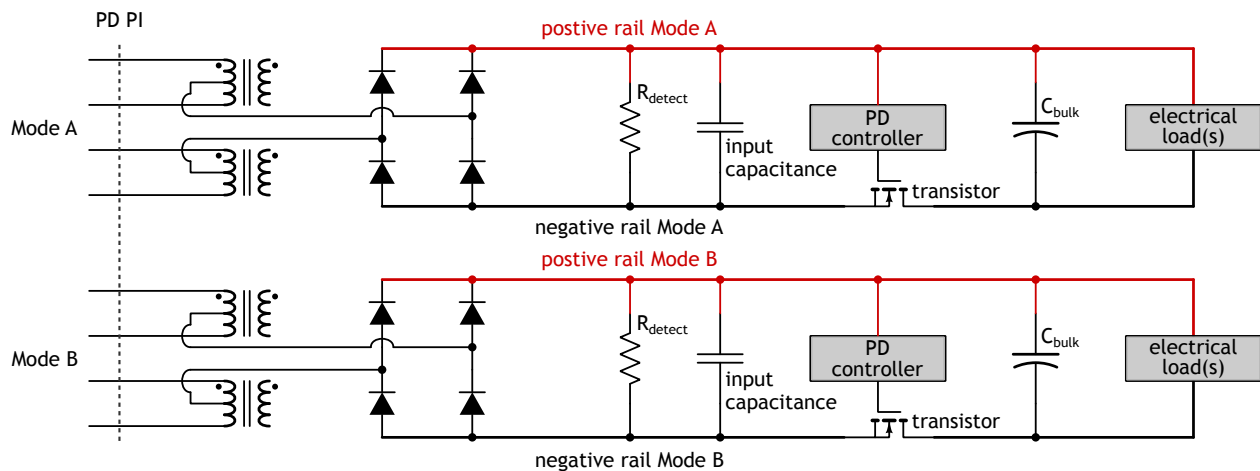
Figure 10: Concept diagram of a dual-signature PD

4.3. Requirements to provide 4-pair power

PSE: 145.2.6.7 4PID requirements

Before a PSE is allowed to provide power to both pairsets, it determines the ability of the PD to accept 4-pair power. Figure 11 shows a decision flow diagram for the PSE. Recall from the detection requirements that power can only be applied to a pairset if that pairset presents a valid detection signature (see Section 3 for detection). Thus, 4-pair power can only be applied when both pairsets present a valid detection signature. This check is the first step in the 4PID determination. All single-signature PDs may be powered over 4-pair. In case of a dual-signature PD, additional checks are required.

START → **Valid detection on both pairsets ?** —Yes→ **PD is single-signature ?** —Yes→ 4-pair permitted

Valid detection on both pairsets ? —No→ 4-pair power not permitted

PD is single-signature ? —No→ **PD is dual-signature ?**

PD is dual-signature ? —No→ 4-pair power not permitted

PD is dual-signature ? —Yes→ **PD is a Type 3/4 dual-signature PD ?** (as discovered through classification on both pairsets) —Yes→ 4-pair permitted

PD is a Type 3/4 dual-signature PD ? —No, or did not check→ **PD shows a valid detection signature on the unpowered pairset when powered in 2-pair?** —Yes→ 4-pair permitted

PD shows a valid detection signature on the unpowered pairset when powered in 2-pair? —No→ 4-pair power not permitted

──► Indicates an abnormal issue or non-compliant PD

Figure 11: Decision diagram to enable 4-pair power

5. Classification

As Type 3 and Type 4 devices are introduced, one of the critical elements of interoperability is the extension of the classification mechanism. Classification is the process by which the PSE and PD negotiate the amount of power the PSE will allocate to the PD. A PD will **request** a certain amount of power, and the PSE will **assign** that power (or less, if it can't meet the demand). This process also allows PSEs and PDs to determine the Type of the device to which they are connected, referred to as mutual identification. This is done in two ways, Physical Layer classification and Data Link Layer (DLL) classification. Classification uses quite a bit of specific terminology, see Section 5.1 for an overview.

Physical Layer classification is performed after detection, but before operating voltage is applied. It consists of a series of class events during which the PSE applies a voltage level to which the PD responds with a certain predetermined current draw (the class signature). Figure 12 shows a timing diagram for a Type 2 PSE powering a Type 2 PD. In this example we see the PSE producing two class events to a Class 4 PD. The PD responds to these two class events by producing a class 4 signature for each event.

This technique has been extended for Type 3 and Type 4 and must be supported by PSEs and PDs. An example of a Type 4 PSE and a Class 7 PD is shown in Figure 13. The PSE produces 5 class events, indicating to the PD that it being assigned to Class 7.

Note that this section describes single-signature classification in detail. Dual-signature classification is substantially different, please refer to the 802.3bt standard for details.

When the PSE applies a voltage in the range of V_{Class} (15.5 V to 20.5 V) for the appropriate amount of time, this is called a class event. A class event is followed by a mark event, this is the application of voltage in the range of V_{Mark} (7 V to 10 V). A mark event allows the PD to recognize that the (previous) class event has concluded.

The relation between the PSE Type and the number of class events that can be produced is listed in Table 4.

The PD's current draw during an individual class event is named the 'class signature'. Five class signatures are defined, numbered 0 through 4, each representing a certain range of current draw

Table 4: PSE Type and corresponding maximum number of class events and assigned Class

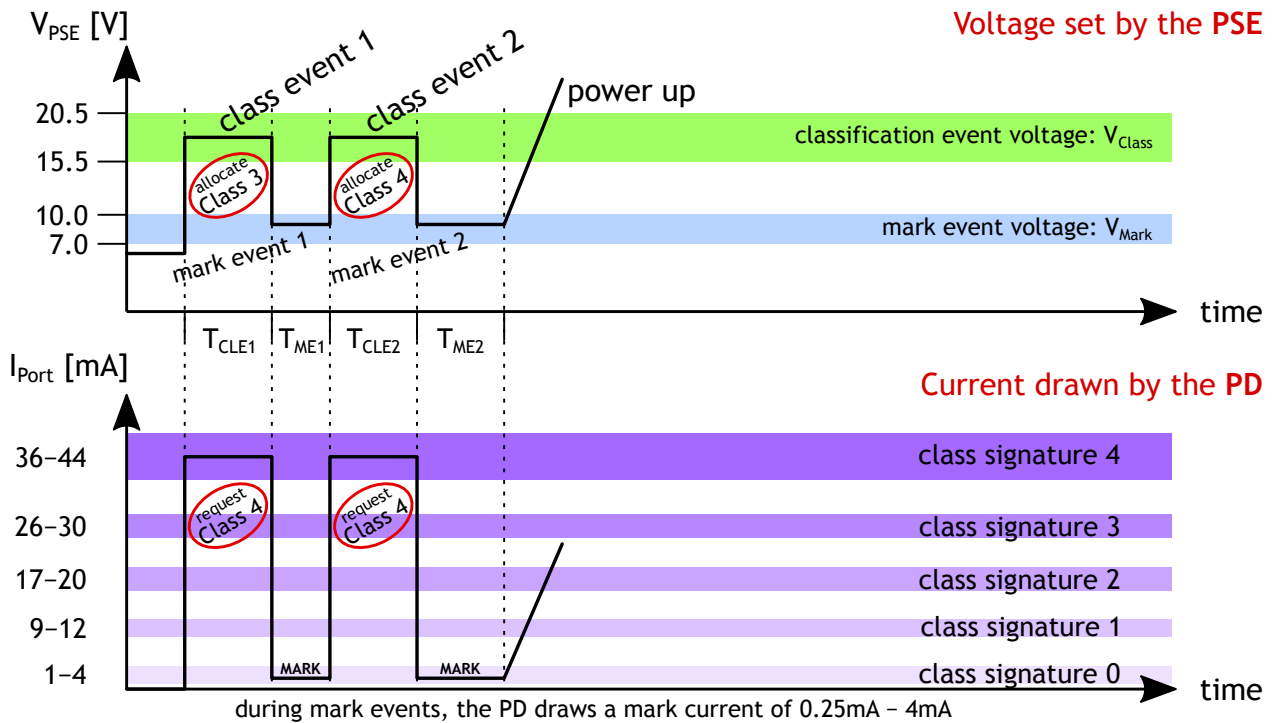| PSE Type	Maximum number of class events	Highest assigned Class
1	1	3
2	2	4
3	4	6
4	5	8

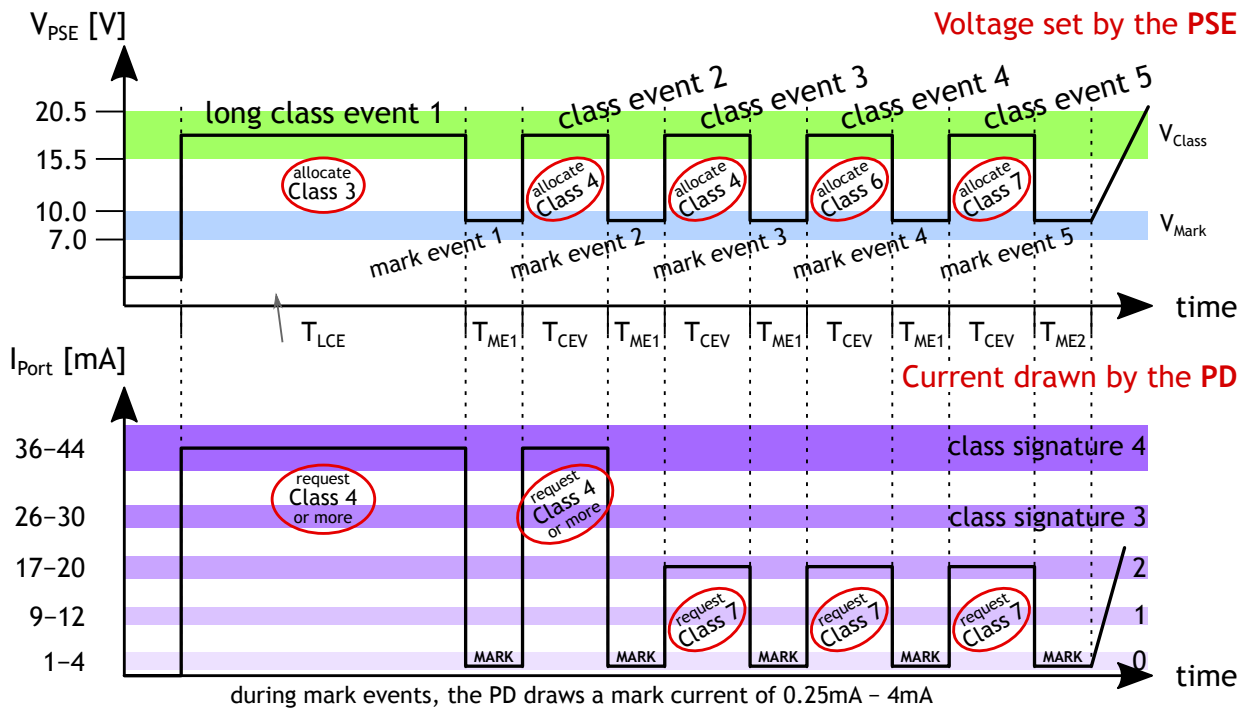Figure 12: Type 2 Physical Layer classification timing diagram

Figure 13: Type 3 and Type 4 Physical Layer classification timing diagram

Table 5: PD requested Class and corresponding class signatures

Requested Class	Class signature produced by the PD				
	Event 1	Event 2	Event 3	Event 4	Event 5
Class 1	1	1	1	/1/	/1/
Class 2	2	2	2	/2/	/2/
Class 3	3	3	3	/3/	/3/
Class 4	4	4	4	/4/	/4/
Class 5	4	4	0	0	/0/
Class 6	4	4	1	1	/1/
Class 7	4	4	2	2	2
Class 8	4	4	3	3	3

This Table shows the class signature the PD produces for each event. PDs that request Class 4 or lower always present the same class signature, regardless of how many events to which the PD is subjected. PDs that request Class 5 through 8, use a different class signature from the third event onward. A class signature denoted /between slashes/ indicates that this number of class events will not occur in a compliant system.

during a class event. The PD's requested Class is defined by the class signatures the PD presents while in the classification voltage range. See Figure 12 and 13 for a graphical representation of the current ranges that the different class signatures represent.

Table 5 shows the class signatures produced by a PD during each class event as a result of its requested Class. The PD signatures for the first two class events are the same as for Type 1 and Type 2 PDs in order to ensure backwards compatibility. During the third class event the PD can change its class signature in order to request a power level greater than Class 4. Only Type 3 and Type 4 PSEs produce more than two class events. This class signature is then repeated for any subsequent class events. The entries noted in slashes (eg. /0/) in Table 5 indicate class events that will not occur in a compliant system, but are there to define behavior for the PD if it is subjected to more class events than expected. To a Type 1 or Type 2 PSE, a Class 5 or higher PD will look like a Class 4 PD. It is only at the third class event that the PD signals that it is requesting more than Class 4 by changing the class signature.

Table 6 shows how assigned Class is determined as a result of PD's requested Class and the number of class events produced by the PSE. Simplifying even further, the assigned Class can be described as the lower of the PD's requested Class and: Class 3 if one class event is produced, Class 4 if two or three class events are produced, Class 6 if four class events are produced, or Class 8 if five class events are produced. Note that the PD class count is incremented on the class to mark transition.

While the PD is requesting power during classification, the PSE is granting or allocating power. This is done through the number of class events that the PSE produces. Table 7 shows the resulting assigned Class based on the PD's requested Class and the number of class events produced by the PSE. It also shows the extension of power demotion that is included in the Type 3 / Type 4 specification.

Table 6: Derivation of assigned Class and power levels for PSE and PD.

Requested Class	Number of class events				
	1	2	3	4	5
Class 1	Class 1	✕	✕	✕	✕
Class 2	Class 2	✕	✕	✕	✕
Class 3	Class 3	✕	✕	✕	✕
Class 4	Class 3	Class 4	Class 4	✕	✕
Class 5	Class 3	Class 4	Class 4	Class 5	✕
Class 6	Class 3	Class 4	Class 4	Class 6	✕
Class 7	Class 3	Class 4	Class 4	Class 6	Class 7
Class 8	Class 3	Class 4	Class 4	Class 6	Class 8

This Table shows the resulting assigned Class, based on the PD's requested Class, and the number of class events that the PSE produced. A PSE may not produce more class events than what is needed to satisfy the PDs power request (eg. a PSE capable of Class 8 power is not permitted to produce five class events when the PD only requests Class 3 — the PSE may only produce one class event in this case).

A green colored cell indicates that the assigned Class is equal to the requested Class, white cells indicate that the PD has been power demoted. A cell marked with a cross indicates that the PSE is not allowed to produce the corresponding number of class events.

The term power demotion refers to a PD being assigned less power than what it asked for. In the existing standard, Type 2 PDs could be demoted to the Type 1 power level when connected to a Type 1 PSE, or power-limited Type 2 PSE, if the PSE only produced one class event. Any PD can be demoted to Class 3 power if the PSE produces only one class event, to Class 4 power if the PSE only produces two or three class events, or to Class 6 power if the PSE only produces four class events. This allows PSEs to supply power to a PD, even if the PSE does not have all of the PD's requested power available, allowing higher power PDs to operate in a reduced mode when connected to lower power PSEs.

As power demotion is critical to the understanding of Physical Layer classification, let's look at an example. Consider a system in which the PSE has 53 W available (sufficient to support Class 5, but not Class 6) and a PD is plugged in. If the PD requests Class 3 (or lower) power, the PSE will give one class event and the PD is allocated its full requested power. If the PD requests up to Class 4 power, the PSE gives either two or three class events and the PD is allocated its full requested power. If the PD requests Class 5 power, the PSE gives four class events and the PD is allocated its full requested power. However, if the PD requests greater than Class 5 power, the PSE only gives three class events and the PD is allocated Class 4 power.

There are a few other changes compared to Physical Layer classification as defined for Type 1 and Type 2 devices:

- Type 3 and Type 4 PSEs produce a long first class event, 88 ms to 105 ms, in order to signal to the connected PD that it is connected to a Type 3 or Type 4 PSE. This allows the PD to use the short MPS timings (see Section 10). In contrast, Type 1 PSEs produce a 6 ms to 75 ms class event, and Type 2 PSEs a 6 ms to 30 ms first class event.

- The maximum length of the second through fifth class events has been reduced to 20 ms in order to reduce the thermal load on the PD controller.

- A PSE is always required to produce a mark event after the final class event prior to the application of operating voltage. This allows PDs to observe the long class event, also if it only requests Class 1 through 3, which results in a single classification event. PDs must also present a valid mark signature during this time. This is different compared to Type 1 and Type 2 classification where a mark event was not required if the PSE only produced a single class event.

- For Type 1 PDs classification is optional. A Type 1 PD that did not implement a classification circuit would present class signature 0 during a class event. A PSE treats such a Class 0 PD as if it requested Class 3. Type 3 PDs are required to support Physical Layer classification, and for this reason Class 0 does not exist for these devices. A Type 3 or Type 4 PSE connected to a Type 1, Class 0 PD, will assign the PD to Class 3, which has an identical power limit.

Table 7: Resulting assigned Class and number of class events based on available PSE power and the requested Class

PSE available power	Requested Class by the PD							
	Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7	Class 8
Class 1	Class 1 (1 EV)	no power	no power	no power	no power	no power	no power	no power
Class 2	Class 1 (1 EV)	Class 2 (1 EV)	no power	no power	no power	no power	no power	no power
Class 3	Class 1 (1 EV)	Class 2 (1 EV)	Class 3 (1 EV)	Class 3 (1 EV)	Class 3 (1 EV)	Class 3 (1 EV)	Class 3 (1 EV)	Class 3 (1 EV)
Class 4	Class 1 (1 EV)	Class 2 (1 EV)	Class 3 (1 EV)	Class 4 (2/3 EV)	Class 4 (2/3 EV)	Class 4 (2/3 EV)	Class 4 (2/3 EV)	Class 4 (2/3 EV)
Class 5	Class 1 (1 EV)	Class 2 (1 EV)	Class 3 (1 EV)	Class 4 (2/3 EV)	Class 5 (4 EV)	Class 4 (2/3 EV)	Class 4 (2/3 EV)	Class 4 (2/3 EV)
Class 6	Class 1 (1 EV)	Class 2 (1 EV)	Class 3 (1 EV)	Class 4 (2/3 EV)	Class 5 (4 EV)	Class 6 (4 EV)	Class 6 (4 EV)	Class 6 (4 EV)
Class 7	Class 1 (1 EV)	Class 2 (1 EV)	Class 3 (1 EV)	Class 4 (2/3 EV)	Class 5 (4 EV)	Class 6 (4 EV)	Class 7 (5 EV)	Class 6 (4 EV)
Class 8	Class 1 (1 EV)	Class 2 (1 EV)	Class 3 (1 EV)	Class 4 (2/3 EV)	Class 5 (4 EV)	Class 6 (4 EV)	Class 7 (5 EV)	Class 8 (5 EV)

5.1. Classification terminology

This section briefly explains the specific terminology that is key to understanding how classification works.

Class: an identifier for the maximum power consumption of a PD. Classes 1 to 8 are defined, representing an increasing maximum amount of power. For example, a Class 4 PD, is a PD that requires up to 25.5 W to operate.

requested Class: the Class that the PD advertises during Physical Layer classification. It represents the amount of power that the PD wants from the PSE. The requested Class is determined by the class signatures that the PD produces when exposed to class events, see Table 5.

assigned Class: the Class that the PSE gives (assigns) to the PD. It represents the amount of power that the PSE allocates for the PD and it also becomes the maximum power limit for the PD. A great number of PSE and PD parameters in the standard depend on the assigned Class. The assigned Class is determined based on the requested Class and the number of class events that the PSE produces, see Table 7.

power demotion: when a PSE assigns a lower Class than what the PD requests. This term is not used in the standard.

class event: a voltage produced by the PSE, in the range of V_{Class}, for a certain amount of time. A class event simultaneously allows the PSE to discover the requested Class of the PD (by measuring what class signature the PD produces), as well as granting a certain amount of power to the PD. Consecutive class events determine what the assigned Class is (see Table 7).

class signature: a current drawn by the PD when the PD voltage is in the range of V_{Class_PD} (i.e. during a class event). Five class signatures are defined, numbered 0 through 4, each corresponding with a certain range of current draw. See Figures 12 and 13 to see the mapping between class current and the corresponding class signature.

class probing: a method for the PSE to determine the PD's requested Class, without the intent to allocate power or power up the PD.

classification reset: is used to reset the PD (make it forget any preceding class events) by reducing the PSE voltage below V_{Reset} for at least T_{Reset}.

mark event: a voltage produced by the PSE, in the range of V_{Mark}, which is less than V_{Class}. A mark event serves to demarcate between class events (or the last class event and power up).

mark event current: the current the PD draws during a mark event, I_{Mark}.

Physical Layer classification: the entire process of a PSE producing class events, a PD exhibiting class signatures, the PSE learning about the requested Class, and finally the PSE assigning a Class to the PD. All this happens before power is applied.

Data Link Layer classification: the process of changing the power allocation after power is applied by negotiating a new maximum power draw using an LLDP TLV (Type/Length/Value) (see 11).

mutual identification: the process devices use to discover each other's Type.

5.2. Class probing

A Type 3 / Type 4 PSE is allowed to perform a 'class probe' in order to discover the requested Class of the PD. To discover this, the PSE may need as many as three class events (in case the class signature of the first class event is 4). The PSE however is not allowed to produce more than a single class event if it has a power budget limited to Class 3 or less.

To enable a PSE with a Class 3 or lower power budget to discover the requested Class, the standard allows the PSE to produce three class events, followed by a classification reset. A classification reset entails the PSE lowering the PD voltage below V_{Reset} (2.8 V) for at least T_{Reset}

(15 ms). This ensures that the PD is completely reset and does not 'remember' the three probing classification events. After the probe is completed, the PSE can then issue a single class event and power up the PD.

During a class probe, the PSE is not required to produce a long first class event, it may use the the 'short' timing of 6 ms to 20 ms in order to speed up the process. The first class event that is part of normal classification is always a long first class event (88 ms to 105 ms).

5.3. Mutual identification

Mutual identification allows a PSE to discover the Type of the PD to which it is connected, and a PD to learn the Type of the PSE to which it is connected. There are a few blind spots in the mechanism, as it is not always possible to determine the connected device Type. However, the specification is designed such that in case of ambiguity, the assigned Class provides all the required information to ensure interoperability.

The PSE can discover the PD's Type based on the requested Class of the PD as shown in Table 8. For PDs that request Class 4 and below it is not possible for the PSE to determine if it's a Type 3 PD, or a Type 1 / Type 2 PD.

Table 8: Deriving PD Type from the requested Class

Requested Class	PD Type
0	Type 1
1, 2, or 3	Type 1 or Type 3
4	Type 2 or Type 3
5 or 6	Type 3
7 or 8	Type 4

For the PD, Type identification of the PSE is sometimes ambiguous, but a distinction between Type 1 / Type 2 PSEs and Type 3 / Type 4 PSEs can always be made by measuring the length of the first class event. A PSE Type identification Table for the PD is provided in Table 9.

An important reason why a PD wants to know the PSE Type is to determine if it can make use of the "short MPS" feature, see Section 10 for details.

Table 9: Deriving PSE Type from the length of the first class event and assigned Class

Duration first class event	Assigned Class	PSE Type
short (≤ 75 ms)	1 – 3	Type 1 or Type 2
	4	Type 2
long (≥ 88 ms)	1 – 6	Type 3 or Type 4
	7 – 8	Type 4

5.4. Data Link Layer Classification

> 79.3.2 Power via MDI TLV
> 145.5 Data Link Layer classification

Data Link Layer (DLL) classification occurs after the PD has received operating power. It operates over the Link Layer Discovery (LLDP) protocol defined in clause 79 of the IEEE 802.3 standard. All PDs, with the exception of Class 0 through Class 3 PDs, are required to support DLL classification. It is optional for PSEs to support DLL classification.

DLL classification is used for a number of reasons. Most commonly, it is used to further refine the power allocated to the PD. For example, a Class 6 PD that requests 51 W through Physical Layer classification can use DLL classification to lower its request to 43 W, allowing the PSE to recover part of the allocated power. DLL classification can also be used by PDs that did not receive their full requested power. For example a Class 8 PD that was assigned to Class 4 (power demotion) can use DLL to request Class 8 power again. If the reason the PSE assigned the PD to Class 4 was due to a lack of available power, it may later be able to reclassify the PD as Class 8 and allocate its full requested power. PDs are not allowed to use DLL to request more power than what they request through Physical Layer classification.

Two fields in the Power over Ethernet TLV (see Section 11) are used to negotiate a new power allocation between a PSE and a PD: the 'PD requested power value' field and the 'PSE allocated power value' field. These two fields are 2 bytes each, and are interpreted as an unsigned integer number.

A PD can request power by putting the amount of power it needs in the 'PD requested power value' field and sending out an updated LLDP frame. The value represents a power level in $1/10^{th}$ of a Watt increments (eg. value 255 represents 25.5 W). The PSE, upon reception, will evaluate the power request, and update the allocation. It will send the new power allocation in the 'PSE allocated power value' field. Only once the PSE and PD are in sync with each other, the power allocation changes. Beware that correct implementation of DLL requires more than this short description!

A third field is noteworthy: the 'PSE maximum available power' field communicates to the PD how much power the PSE has available. It is *not* an allocation, rather it is a hint of what kind of power request the PSE would be willing to grant. To actually get the power, the PD must always use the 'PD requested power value' field and observe what happens to the 'PSE allocated power value' field.

After a DLL transaction is successfully completed, both the allocated power (PSEAllocatedPowerValue), and the PD power limit (PDMaxPowerValue) will be set to the new negotiated power level. As a result of this, it is possible that the assigned Class changes! For example, a PD that request Class 6, but gets power demoted to Class 4, will initially start with PSEAllocatedPowerValue and PDMaxPowerValue at a value of 255 (representing Class 4 power limit of 25.5 W). Later on, the PD requests what it really needs (50 W) and this is granted by the PSE, provided that it is able to allocate the requested amount of power. Now PSEAllocatedPowerValue and PDMaxPowerValue are both equal to 500. The PD has now been re-assigned to Class 6, which means that all of the class dependent requirements may have changed. A mapping between the

values of PSEAllocatedPowerValue and PDMaxPowerValue and the resulting assigned Class are listed in Table 10.

Table 10: Relation between negotiated DLL power and the assigned Class

PSEAllocatedPowerValue (PSE) PDMaxPowerValue (PD)	Assigned Class
1 – 39	Class 1
40 – 65	Class 2
66 – 130	Class 3
131 – 255	Class 4
256 – 400	Class 5
401 – 510	Class 6
511 – 620	Class 7
621 – 999	Class 8

6.   Autoclass

> PSE: 145.2.8 PSE classification of PDs and mutual identification
> 145.2.8.2 Autoclass (optional)
>
> PD: 145.3.6.2 Autoclass (optional)
> 145.3.8.2 Input average power

Included in the 802.3bt standard is an optional extension of Physical Layer classification, named Autoclass. Autoclass allows the allocation of PSE power supply budget to the PD to be far better optimized than what is possible using 'normal' Physical Layer classification, or even Data Link Layer classification.

A PD can only correctly use Autoclass if it is able to put itself in a state where it consumes the maximum amount of power it needs.

The concept of Autoclass is that a PD can request "Autoclass" during Physical Layer classification. This is done by transitioning its given (non-zero) class signature to a class signature of zero after approximately 81 ms during the first classification event. Apart from this, classification is as outlined earlier. See Figure 14 for an overview of the Autoclass Physical Layer classification specifics.

After the system has been powered, the PD will draw the maximum amount of power it needs. During this time, the PSE may measure the power consumption. That measurement of power will reveal to the PSE precisely how much power it needs to allocate for this particular PD. Because it is based on a measurement of the actual system, the resulting power budget is compensated for the actual power requirement of the PD in use, the actual cable power losses, and the power measurement uncertainty of the PSE in use.
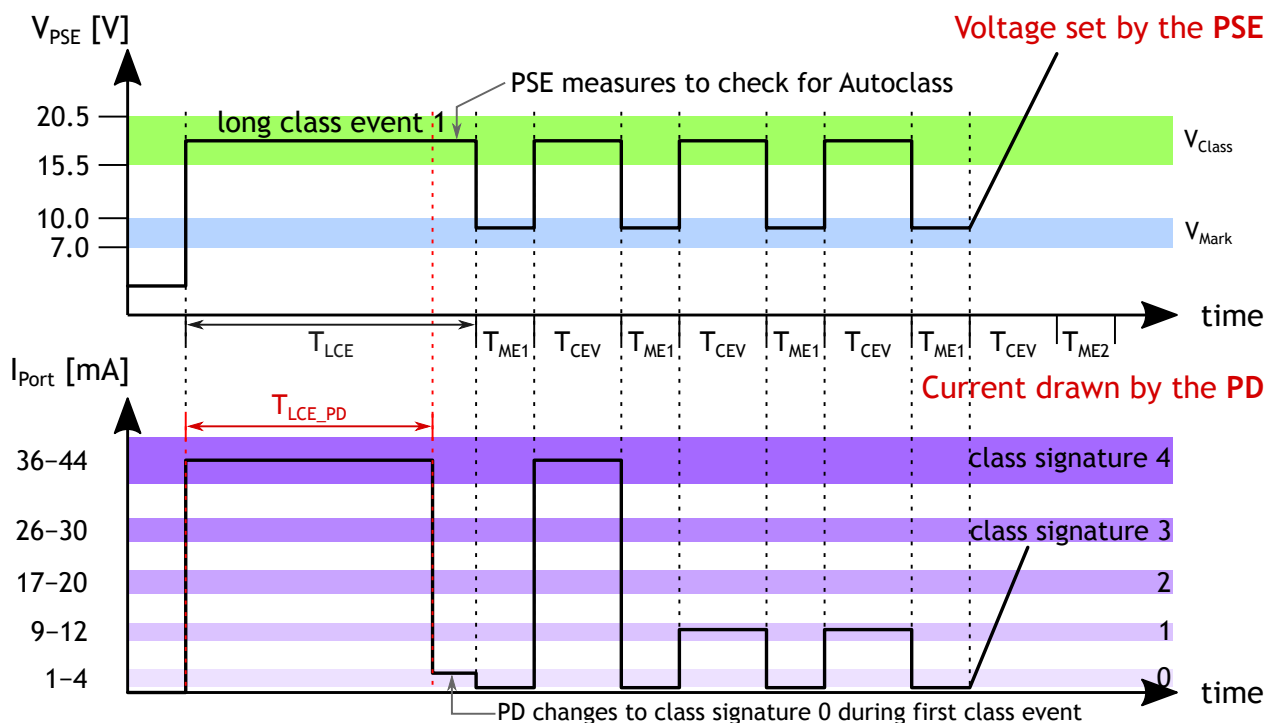
Figure 14: Autoclass Physical Layer classification timing overview

To put this in numbers, let's say we have a Class 8 PD that requires 65 W. In our example, it is connected to the PSE by a 25 m patch cord of AWG 23. The PD is connected and gets assigned to Class 8. The PSE power supply is at 55 V. Without Autoclass, most PSEs would allocate 90 W, since that is the worst-case amount of power the PSE could be expected to source.

With Autoclass, the PD consumes 65 W (at the PD PI) after startup. The PSE power measurement will reveal that it is sourcing 66.5 W. For Class 8, a PSE is required to increase this with 1.25 W of margin, resulting in an allocation of 67.8 W. That is a difference of 22.2 W, a saving of almost 25% compared to the worst-case allocation.

Note that without Autoclass, the PSE has a number of options to improve the power allocation; using DLL classification it can learn the actual maximum power consumption of the PD, if the PSE has a better than minimum voltage power supply it can include that in it's budget calculation, and if the PSE is aware of the actual cable resistance, it can compensate its allocation for this as well. The PSE could for instance use the newly defined LLDP PoE measurements to discover the DC resistance of the cabling.

Support for Autoclass for a PSE is optional. In the case that the PSE does not implement it, but the PD does request Autoclass, the procedure outlined above still occurs for the PD, but the PSE will not make the Autoclass measurement, nor will it adjust it's power budget. The power budget will be based on the Class assigned by the PSE.

Apart from Physical Layer Autoclass, it is also possible to request an Autoclass measurement from the PSE by using the LLDP protocol. This allows a PD to use Autoclass, even if it is unable to meet the Physical Layer Autoclass timing requirement to enable the PD maximum power draw immediately after power up.
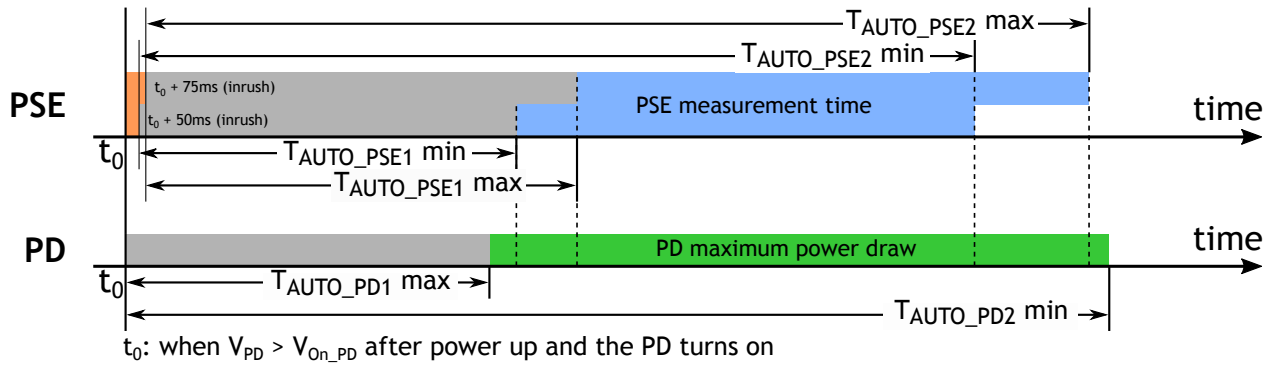
Figure 15: Autoclass PSE power measurement and PD maximum power draw timing

A common reference time t_0 is used to show both the PSE and PD Autoclass timings from a common reference. The orange part of the graph is the PSEs minimum and maximum inrush time. The grey area is the wait time for the PSE before starting the Autoclass measurement. Blue is the PSE measurement time. Both minimum time periods and maximum time periods are shown, showing the extreme points of the measurements.

For the PD, grey represents the time it has to transition the load to maximum power draw. Throughout the green period, the PD is required to draw the maximum power.

6.1. PSE Autoclass requirements

Nearly immediately after power has been applied, the PD must consume the highest amount of power it will need for about 1.5 s and the PSE will measure the sourced power at the PSE PI to establish the power budget. A small delay is incorporated to allow the PD to conclude inrush and activate its load. The timings are shown in Figure 15 and Table 11. A PSE, after V_{PSE} exceeds 30 V, will wait for T_{AUTO_PSE1}, and then measure power until T_{AUTO_PSE2}. During the window between T_{AUTO_PSE1} and T_{AUTO_PSE2}, the PSE measures continuously and averages the power consumption with a sliding window of 150 ms to 300 ms. This results in a series of power measurements, the highest one of these determines $P_{Autoclass}$. The minimum power allocation for an Autoclass PSE is Class 1 power.

The PSE is required to add a small amount of margin to the power budget (P_{ac_margin}), with the purpose of accommodating for increased cable resistance should the cable heat up during operation. This extra margin is 0.5 W for assigned Class 1 through 4, 0.75 W for assigned Class 5 and Class 6, and 1.75 W for Class 7 and Class 8.

A potential caveat exists for PSEs that have assigned Class 1 through 4, as these are permitted to switch between 2-pair and 4-pair mode at any time. These PSEs should take into account that if they measure Autoclass in 4-pair mode, they will need to allocate additional power when the PSE switches to 2-pair mode. This extra allocation, only required while operating in 2-pair mode, is called P_{ac_extra}, and is calculated with the highest resistance channel:

$$P_{ac\_extra} = \left(\frac{P_{Autoclass}}{V_{Port\_PSE-2P} \ min} \right)^2 \times \frac{R_{Ch}}{2}$$

Alternatively, such a PSE can choose to make the Autoclass measurement in 2-pair mode, so that

it is guaranteed to have measured the worst-case power consumption. This results in a more optimal allocation compared to making use of P_{ac_extra}.

6.2. PD Autoclass requirements

After V_{PD} exceeds V_{On_PD} (30 V to 42 V), the PD has a maximum time of T_{AUTO_PD1} to reach a state where it is consuming it's maximum power. This is not a lot of time, nor does it allow any kind of thermal settling. Depending on the PD application it may be necessary to intentionally draw a little bit more power to accommodate thermal effects during operation. Note that Autoclass works side by side with normal classification: the PD must still meet all of the regular Physical Layer classification requirements. Autoclass can only serve to give power back to the PSE. It is also possible that an Autoclass PD gets power demoted to a lower Class, in this case the PD is still required to meet the power limitations of that lower Class.

The PSE time parameters, T_{AUTO_PSE1} and T_{AUTO_PSE2}, are referenced from the **end** of inrush, compared to the start of the T_{AUTO_PD1} timer which is at the **beginning** of inrush. There is thus a 50 ms to 75 ms difference in the starting point of the PSE and PD Autoclass timing parameters. The PSE waits for T_{AUTO_PSE1} before starting the measurement, and may measure for up to T_{AUTO_PSE2} (referenced from the end of PSE inrush). The period during which the PSE measures the power sourced is completely within the period where the PD is required to draw maximum power, as shown in Figure 15.

Table 11: Autoclass power measurement and maximum power draw timing

Description	Symbol	min	max	
PD maximum power draw start time	T_{AUTO_PD1}		1.35	s
PD maximum power draw end time	T_{AUTO_PD2}	3.65		s
PSE power measurement start time	T_{AUTO_PSE1}	1.4	1.6	s
PSE power measurement stop time	T_{AUTO_PSE2}	3.1	3.5	s

7. Inrush

> PSE: 145.2.10.6 Output current during power up
> PD: 145.3.8.3 Input inrush current

The inrush phase is the last phase before operating power is applied and the system begins normal operation. Inrush is the controlled application of power, to gently switch over from a non-powered state to a powered state. A large portion of interoperability problems stem from PD implementations that do not correctly implement inrush requirements. Designers are encouraged to carefully read this section, as well as the relevant sections on inrush in the 802.3bt standard.
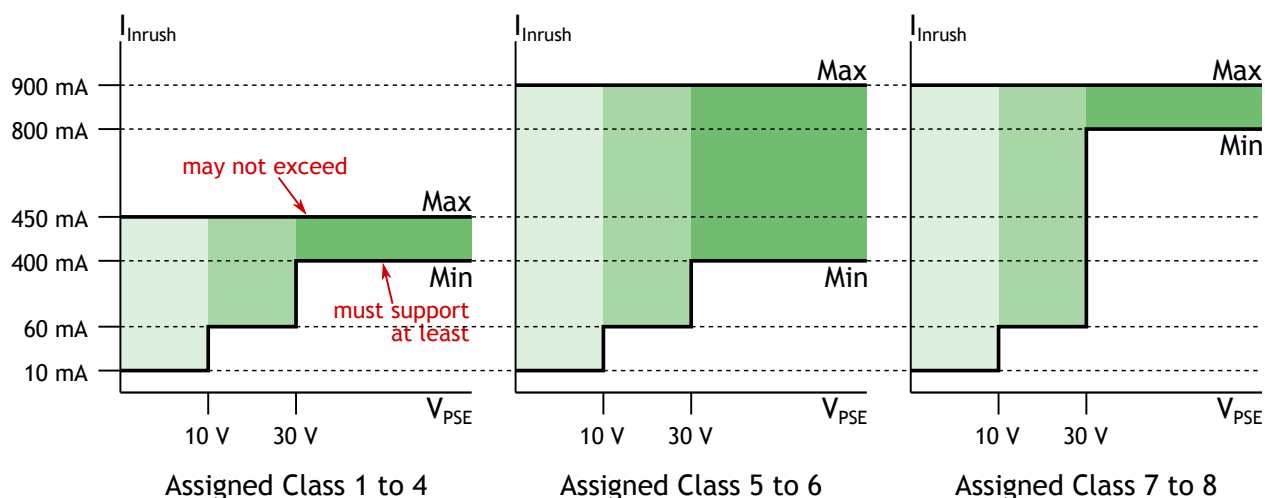
Figure 16: PSE total inrush current, I_{Inrush}

7.1. PSE inrush

After classification, if the PSE decides to power the PD, it will first go through the inrush phase. During inrush the PSE limits the amount of current being delivered for at least 50 ms and up to 75 ms. The PSE current limit for inrush (I_{Inrush}) depends on the assigned Class as well as the voltage at the PSE PI. This is shown in Figure 16. In order to decrease the thermal load on the PSE current limiter, the full inrush current is only required when the voltage has risen above 30 V.

In addition to the requirement on total inrush current (I_{Inrush}), the PSE limits the current on a pairset to $I_{Inrush-2P}$. For assigned Class 1 to 4, $I_{Inrush-2P}$ is equal to I_{Inrush}, at 450 mA. For assigned Class 5 to 8, $I_{Inrush-2P}$ is 600 mA. The consequence of this is that PSEs can use 2-pair or 4-pair inrush for Class 1 to 6, but must inrush over 4-pair for Class 7 or 8 in order to meet both the pairset current limit and the total minimum inrush current.

7.2. PD inrush

For the PD, the inrush phase begins when the PD voltage (V_{PD}) exceeds the PD turn on voltage (V_{On_PD}) which is anywhere in the range of 30 V to 42 V. This moment is named t_0. The inrush phase ends 80 ms after t_0.

During the first 50 ms of the inrush phase, the PSE limits the inrush current to a total current of I_{Inrush}, and limits each pairset to a current of $I_{Inrush-2P}$. The values of these input currents depend on the assigned Class and on the voltage at the PSE PI (see Figure 16). During the "PSE inrush" phase there is no particular requirement on the PD.

After 50 ms the PD is in the power delay phase, which lasts for 30 ms. Somewhere in the PD power delay phase, the PSE switches from its inrush state, to its powering state. During this time, the PD needs to meet the following requirements:

1. Have a total input **current** less than I_{Inrush_PD} (parameter depends on the assigned Class) and have a pairset current less than $I_{Inrush\text{-}2P}$ on each pairset.

2. Have an input **power** less than Class 3 (13 W) or the PD's requested Class, whichever is lower.

The reason for item 2 above is that the PSE can switch out of inrush at any time between 50 ms to 75 ms and immediately enforce the assigned Class.

To meet these requirements, it is highly recommended to keep the PD electrical load disabled until the PD is solidly into the 'power applied' state, as shown in Figure 17. If the load is turned on prematurely, part (or all) of the current the PSE provides is taken by that load, and prevents the bulk capacitor from charging.

There are several methods employed by PDs to complete inrush. Ultimately, only one of these can guarantee the PD to successfully complete inrush and achieve interoperability with all PSEs.

1. The PD does not perform any form of inrush control. Once the PSE turns on and the PD voltage has risen above V_{On_PD}, the PDs power switch turns on and stays on. This has the effect that the voltage at the PD PI (and PSE PI) collapses to near zero. As can be seen in Figure 16, PSEs are allowed to deliver small currents at this voltage. Unless the PSE provides significantly more inrush current that the minimum required by the standard, inrush will fail because the bulk capacitor of the PD will not be charged when the PSE transitions from inrush to power on. Such a PD fails to comply with the V_{Off_PD} requirement, which states that below 30 V a PD must be "turned off".

2. The PD does not perform inrush current limitation, but adheres to the V_{On_PD}/ V_{Off_PD} requirement. Whenever the PD turns on, it will transfer the charge in its small input capacitor (maximum 120 nF) through the power switch to its larger bulk capacitor. Once the voltage is reduced to V_{Off_PD}, the power switch turns off. This will cause the power switch to rapidly turn on and off, essentially acting as a switched-capacitor current regulator. Besides being inefficient, such a PD violates the requirement that a PD may not oscillate during turn on.

3. The PD limits the inrush current below I_{Inrush_PD}. Once the PD turns on, it limits the inrush current to a value below I_{Inrush_PD} (400 mA or 800 mA). The voltage at the PD and PSE PI does not collapse back down, but quickly rises to the nominal operating voltage. Because the PD controls the inrush current, it's power switch carries the thermal load of inrush, rather than the PSE. This is in the only compliant and interoperable way a PD can charge its bulk capacitance.

   Note that at t_0 + 50 ms the PD, in addition to drawing less than I_{Inrush_PD}, must conform to the power limits of the "Power delay" state. This power limit is 13 W if the PD requested Class 3 or higher, or 3.84 W or 6.49 W if the PD requested Class 1 or Class 2 respectively. The corresponding input current limits are significantly below the I_{Inrush_PD} value.

After the power delay state, the PD has reached the nominal operating phase and can draw power up to the assigned Class. Refer to Section 8 for more information.
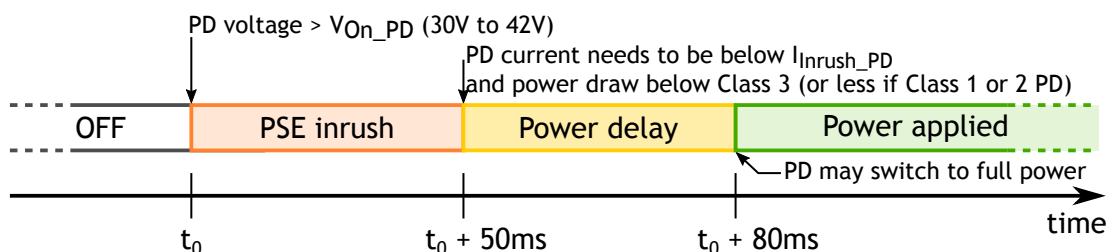
PD voltage > V_{On_PD} (30V to 42V)

PD current needs to be below I_{Inrush_PD}
and power draw below Class 3 (or less if Class 1 or 2 PD)

| OFF | PSE inrush | Power delay | Power applied |

PD may switch to full power

time

t_0 t_0 + 50ms t_0 + 80ms

Figure 17: PD inrush time phases

8. Operating power

PSE: 145.2.8 PSE classification of PDs and mutual identification
145.2.10 Power supply output
145.2.11 Power supply allocation

PD: 145.3.6 PD classification
145.3.8 PD power
145.3.8.2 Input average power

Following successful detection and classification a system usually spends most of its time in the nominal powering state. During this time the PSE checks for abnormal conditions, like overloads or short-circuits, as well as checking that the PD remains connected. As a result of classification, the PD will either have been assigned to the Class it requested, or may have been assigned to a lower Class (referred to as power demotion). In either case, both the PSE and the PD are required to conform to the requirements of the assigned Class. The assigned Class can also change after the system is powered on, as a result of Data Link Layer classification (see 5.4).

8.1. PD power limits

There are three possible power limits that may apply to a PD. Table 12 indicates under which conditions the various power limits apply.

1. P_{Class_PD} is the maximum amount of input average power for a given Class at the PD PI. When a PD is powered, and has not requested Autoclass or performed DLL classification, the power limit for the PD is P_{Class_PD} for the assigned Class.

2. **PDMaxPowerValue** is a maximum power value expressed in 1/10th of a Watt. This value is the result from Data Link Layer Classification, which allows the PD and PSE to negotiate the amount of allocated power. After a successful DLL negotiation, PDMaxPowerValue is the applicable limit. The corresponding PDMaxPowerValue is lower than or equal to P_{Class_PD}. See Section 5.4 for details on DLL classification.

3. $P_{Autoclass\_PD}$ is the amount of power the PD drew during Physical Layer Autoclass, or during Data Link Layer Autoclass. A PD that requested Autoclass during Physical Layer classification is required to consume less than $P_{Autoclass\_PD}$ (which is the power it drew right after

Table 12: Applicable power limit for a PD

Condition	Applicable limit
Power applied after classification (without Autoclass)	P_{Class_PD}
Power applied after classification (with Autoclass)	$P_{Autoclass\_PD}$
After the successful conclusion of a DLL power negotiation	PDMaxPowerValue
After the successful conclusion of a DLL Autoclass request	$P_{Autoclass\_PD}$

Table 13: Derivation of assigned Class and power levels for PSE and PD.

Assigned Class	PD limits		PSE worst case	
	P_{Class_PD} (W)	P_{Peak_PD} (W)	P_{Class} (W)	P_{Peak} (W)
1	3.84	5.00	4.00	5.47
2	6.49	8.36	6.70	8.87
3	13.00	14.40	14.00	16.07
4	25.50	28.30	30.00	34.12
5	40.00	42.00	45.00	47.68
6	51.00	53.50	60.00	63.62
7	62.00	65.10	75.00	79.83
8	71.30	74.90	90.00	96.36

being powered; see Section 6). $P_{Autoclass\_PD}$ is also the applicable power limit after a PD performs DLL Autoclass. The value of $P_{Autoclass\_PD}$ is always lower than P_{Class_PD} for the assigned Class.

The available amount of input power after Physical Layer Classification, P_{Class_PD}, is shown in Table 13, as well as the amount of peak power that can be drawn by the PD. The two rightmost columns show the power that the PSE must be able to source in order to support the PD under worst-case conditions. Worst case conditions include the PD drawing the maximum amount of power allowed, the PSE's supply voltage equaling the minimum allowed, and the link section connecting the two devices having the maximum amount of resistance permitted.

The PD needs to take several requirements on input power into account. First, there is a limit on peak power that the PD is not allowed to exceed. In the standard this is referred to as P_{Peak_PD}; the values depend on the assigned Class and are listed in Table 13. Also, within a 1 second sliding window, the PD may exceed P_{Class_PD}, PDMaxPowerValue, or $P_{Autoclass\_PD}$ (whichever is applicable), for up to 50 ms, up to P_{Peak_PD}. Finally, the average input power over a 1 second sliding window may not exceed P_{Class_PD}, PDMaxPowerValue, or $P_{Autoclass\_PD}$ (whichever is applicable). This includes any peak power that is drawn inside of that window. An example is shown in Figure 18, the input average power for the example window includes all of the green shaded areas.
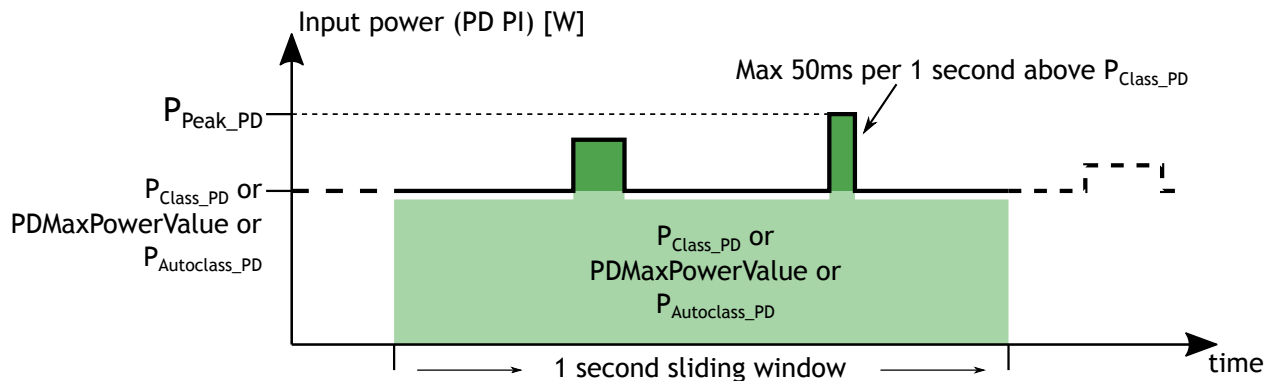
Figure 18: PD input average power and peak power

8.2. PSE power

> 145.2.8.5 Continuous output current capability in the POWER_ON state
> 145.2.8.7 Overload current
> 145.2.8.8 Output current − at short circuit condition

Since the original 802.3af standard, the PSE's output power requirements have been specified using an output current. In conjunction with the permitted output voltage at the PI, V_{Port_PSE-2P}, this effectively becomes an output power requirement. Most of the specification in Clause 145, deals with behavior on a per pairset basis. This, combined with specifying output current as an indirect way to specify power, creates a fairly complicated specification.

The PSE has both a minimum amount of power that it is required to support for a given assigned Class, named the "PSE lowerbound template", as well as a maximum current/power output limit, named the "PSE upperbound template". The PSE is required to prevent current to exceed the upperbound template. The templates are different for Type 3 and for Type 4 PSEs. The operating template for Type 3 is shown in Figure 19 and for Type 4 in Figure 20.

The purpose of the lowerbound template is to require the PSE to support a compliant PD and is specified such that it guarantees a power level of P_{Class_PD} at the PD PI, and a peak power of P_{Peak_PD} at the PD PI. There is no mandatory margin, a PSE is permitted to employ optimization of the supplied current, as long as it can guarantee P_{Class_PD} and P_{Peak_PD}.

The purpose of the upperbound template is to protect the PSE and the cabling infrastructure from over current. The requirement on the PSE is to remove power from the pairset before the current exceeds the upperbound template. In the remainder of this subsection, we'll describe the different segments of the upperbound and lowerbound templates.

The lowerbound template for the 'normal operating current' region (the rightmost region) has a value of I_{Con-2P}, the operating current the PSE supports indefinitely. When the PSE is powering the PD in 2-pair mode, I_{Con-2P} is defined as follows:

$$I_{Con-2P} = \frac{P_{Class}}{V_{PSE}} \qquad (1)$$

$I_{Con\text{-}2P}$ The current the PSE supports on a pairset indefinitely

P_{Class} The amount of output power required to supply P_{Class_PD} at the PD PI, given a certain cable resistance (R_{Chan}) and the actual PSE voltage (V_{PSE})

V_{PSE} The actual PSE PI voltage

In other words, for the 2-pair case, $I_{Con\text{-}2P}$ is simply the current that is needed, for the given system, to provide P_{Class_PD} at the PD PI.

If the PSE is providing power in 4-pair, it needs to meet the same requirement, however since everything is specified on a per pairset basis, it needs to convey a total 4-pair power requirement by specifying two different pairset currents. In addition, as is explained in Section 9, a PSE is only required to support up to $I_{Con\text{-}2P\text{-}unb}$ on a pairset. For 4-pair, $I_{Con\text{-}2P}$ is defined as:

$$I_{Con\text{-}2P} = \min\left(I_{Con} - I_{Port\text{-}2P\text{-}other}, I_{Con\text{-}2P\text{-}unb}\right) \tag{2}$$

$I_{Con\text{-}2P}$ The current the PSE supports on a pairset indefinitely

I_{Con} The total amount of output current the PSE supports (over both pairsets) to provide P_{Class_PD} at the PD PI

$I_{Port\text{-}2P\text{-}other}$ is the actual current flowing on the 'other' pairset

$I_{Con\text{-}2P\text{-}unb}$ is the minimum current a PSE must support on a pairset to cope with pair-to-pair current unbalance

The standard defines $I_{Port\text{-}2P}$ as the current on a given pairset (either Alternative A or Alternative B), with $I_{Port\text{-}2P\text{-}other}$ being the 'other' pairset current. For example, if we evaluate the $I_{Con\text{-}2P}$ equation for Alternative B, then $I_{Port\text{-}2P}$ is the current on Alternative B, and $I_{Port\text{-}2P\text{-}other}$ is the current on Alternative A. This notation allows a single equation to apply to both pairsets. A different way of specifying $I_{Con\text{-}2P}$ would be:

$$I_{Con\text{-}2P} \text{ for Alternative A} = \min\left(I_{Con} - I_{Port\text{-}2P} \text{ (Alternative B)}, I_{Con\text{-}2P\text{-}unb}\right)$$
$$I_{Con\text{-}2P} \text{ for Alternative B} = \min\left(I_{Con} - I_{Port\text{-}2P} \text{ (Alternative A)}, I_{Con\text{-}2P\text{-}unb}\right) \tag{3}$$

The min() function ensures that the required current on any given pairset is never higher than $I_{Con\text{-}2P\text{-}unb}$. The other part, $I_{Con} - I_{Port\text{-}2P\text{-}other}$, results in the amount of current that is required to support a total current of I_{Con} over both pairsets. The total current I_{Con} is defined as:

$$I_{Con} = \frac{P_{Class}}{V_{PSE}} \tag{4}$$

I_{Con} The total amount of output current the PSE supports (over both pairsets) to provide P_{Class_PD} at the PD PI

P_{Class} The amount of output power required to supply P_{Class_PD} at the PD PI, given a certain cable resistance (R_{Chan}) and the actual PSE voltage (V_{PSE})

V_{PSE} The actual PSE PI voltage

Finally P_{Class}, the amount of power the PSE needs to support to provide P_{Class_PD} at the PD PI, is defined as:

$$P_{Class} = V_{PSE} \times \left(\frac{V_{PSE} - \sqrt{V_{PSE}^2 - 4 \times R_{Chan} \times P_{Class\_PD}}}{2 \times R_{Chan}}\right) \tag{5}$$

P_{Class} The amount of output power required to supply P_{Class_PD} at the PD PI, given a certain cable resistance (R_{Chan}) and the actual PSE voltage (V_{PSE})

V_{PSE} The actual PSE PI voltage

R_{Chan} The actual loop resistance between the PSE and the PD

P_{Class_PD} The maximum average amount of power a PD may draw for a particular assigned Class

Because the P_{Class} equation depends on the **actual** cable resistance (not the worst-case) and the actual PSE output voltage, this allows the PSE to optimize the power it budgets for the PD.

The next segment to the left of the lowerbound template serves to allow a PD to draw peak power. The requirements for peak current are identical to those of the normal operating current, with the exception that P_{Peak} is determined using P_{Peak_PD}, and that the permitted unbalance is equal to $I_{Peak-2P-unb}$. A PSE must support this output current for at least T_{CUT} min (50 ms) in any 1 second window.

The final segment (leftmost) of the lowerbound template is the transient current. For a time duration of at least T_{LIM-2P} min, the PSE is required to actively limit the output current to any value between I_{LIM-2P} and the upperbound template. For a Type 3 PSE this timer period is 10 ms, for a Type 4 PSE it is 6 ms. Therefore, during this time it is impossible for a PD to violate the upperbound template.

The upperbound template for Type 3 PSEs consists of fixed values, see Figure 19. The upperbound template for Type 4 PSEs is slightly more complex (see Figure 20). A Type 4 PSE may source up to 1.75 A for up to T_{CUT} max (75 ms), and source up to 1.3 A for up to 4 seconds continuously. In order to comply with IEC62386-1 and IEC60950 requirements for Limited Power Sources, the standard disallows a PSE to source more than 99.9 W for longer than 4 seconds. This is done through the parameter I_{LPS-2P} which is part of the Type 4 PSE upperbound template.

I_{LPS-2P} is equal to 0.85 A when the PSE is operating in 2-pair mode (this current level matches with the Type 3 PSE upperbound template), and is defined in 4-pair mode as:

$$I_{LPS-2P} = \min \left(\frac{P_{Type}\ max}{V_{PSE}} - I_{Port-2P-other}, 1.3 \right) \tag{6}$$

I_{LPS-2P} The maximum amount of current a Type 4 PSE is allowed to source on a pairset

$P_{Type}\ max$ The maximum amount of power a PSE may source under nominal conditions

V_{PSE} The actual PSE PI voltage

$I_{Port-2P-other}$ is the actual current flowing on the 'other' pairset.

8.3. Powering configurations

PSE: 145.2.4 PSE PI
PD: 145.3.2 PD PI

Depending on the PSE Type, different mappings of power supply voltage rails to connector pins are allowed, resulting in different output polarities at the PSE PI. At all time the two conductors of a pair will be at the same polarity. The permissible configurations, along with the pin mapping

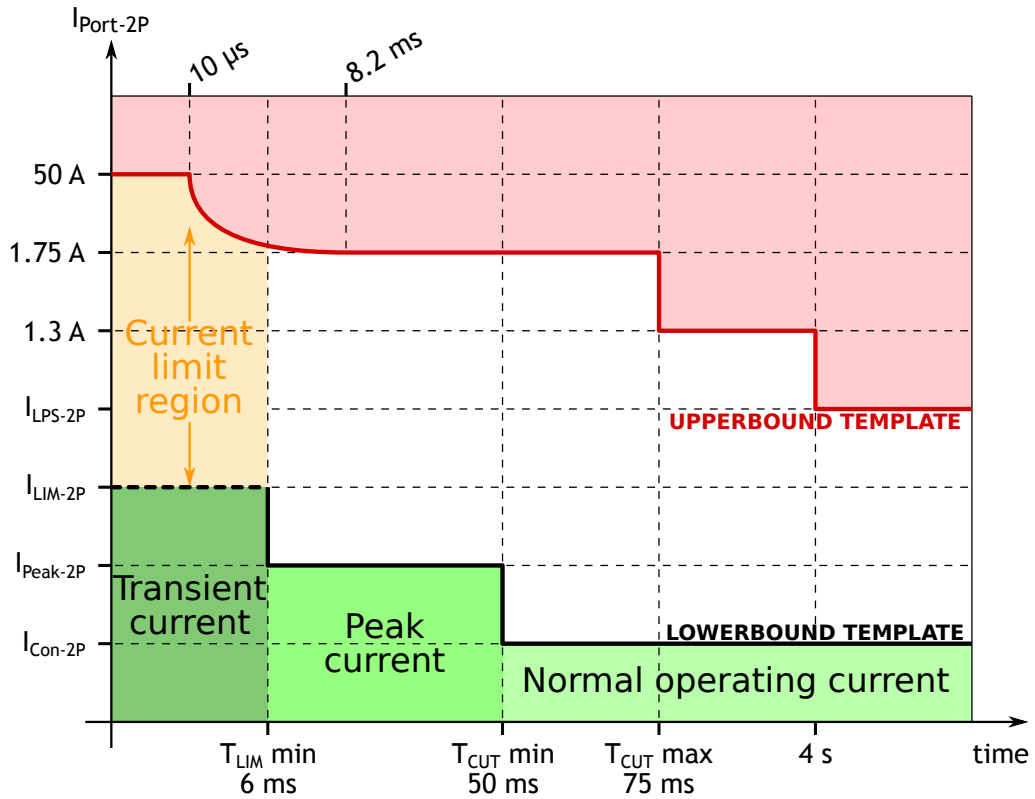Figure 19: Type 3 PSE operating current template

Figure 20: Type 4 PSE operating current template

are shown in Figure 21. When operating in 2-pair mode, either Alternative A will be powered, or Alternative B. Type 1 and Type 2 PSEs exclusively power in 2-pair mode, and may choose any of the configurations listed as "Type 1 & Type 2" in Figure 21.

Type 3 PSEs are permitted to use any of the four possible polarity permutations, while Type 4 PSEs are only permitted to use one polarity configuration, as shown in in Figure 21. The chosen polarity configuration by the PSE is the same during all states (detection, classification, powering).

PDs, of any Type, are required to support any polarity configuration. This means a form of rectification is needed in the PD, either two full diode bridges, or an active rectifier that makes use of transistors. Type 3 and Type 4 PDs are furthermore required to support both 2-pair and 4-pair input power.

A very common implementation of 4-pair PSEs is to directly tie the positive pairs to the power supply and implement a switch and current measurement on each of the negative pairs. When the PSE is operating in "2-pair mode", positive current flows through **two** positive pairs, whereas the return current flows through **one** negative pair back. The standard still considers this "2-pair mode", even though power is actually delivered through 3-pairs.

While a Type 1 and Type 2 PSE can only make use of 2-pair, Type 3 and Type 4 PSEs that have assigned Class 1 through Class 4 to a PD, are permitted to use either 2-pair or 4-pair powering. The PSE may also freely switch between these two modes while providing power. When the PSE provides power in 2-pair, it is always on the same Alternative. The choice of the 2-pair Alternative can only change after a power cycle on the port.

If the PSE has assigned Class 5 or higher, it must provide power in 4-pair mode exclusively. See Table 14 for an overview of what options exist pertaining to 2-pair and 4-pair power for the different PSE Types and assigned Classes.
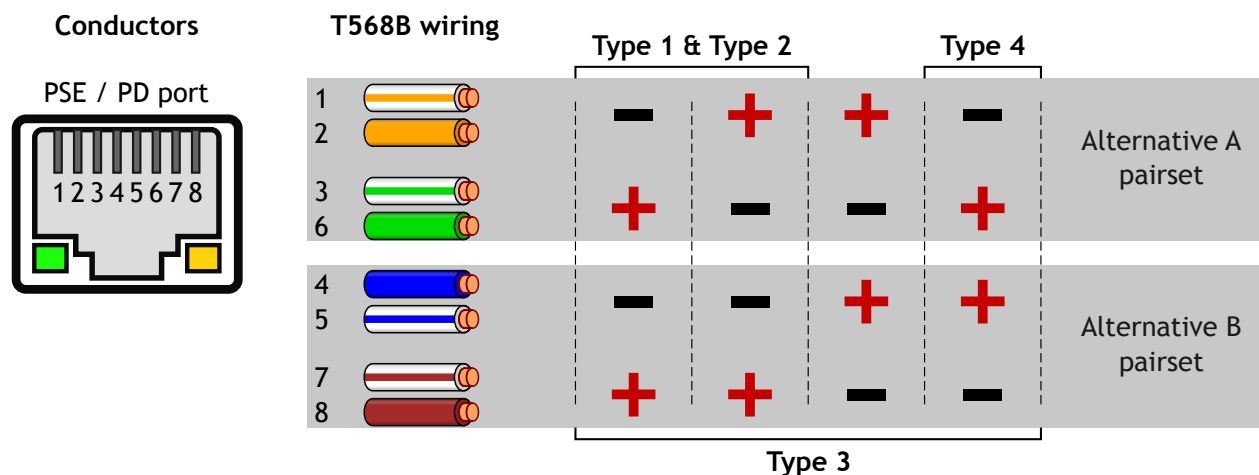

Figure 21: PSE pinout configuration and permissible power supply polarity

Table 14: 2-pair and 4-pair powering rules

	Assigned Class							
	1	2	3	4	5	6	7	8
Type 1 PSE	2P	2P	2P					
Type 2 PSE	2P	2P	2P	2P				
Type 3 PSE	2P 4P	2P 4P	2P 4P	2P 4P	4P	4P		
Type 4 PSE	2P 4P	2P 4P	2P 4P	2P 4P	4P	4P	4P	4P

9. Current unbalance

> PSE: 145.2.10.5.1 PSE pair-to-pair current unbalance
> PD: 145.3.8.9 PD pair-to-pair current unbalance
> Annex with more unbalance information: Annex 145A

When a system is operating in 4-pair mode, two pairs are at the positive voltage and carry current toward the PD, and two pairs are at the negative voltage and carry that same total current back to the PSE. The two pairs (either negative or positive) are electrically parallel to each other, and included in this parallel path are a number of components inside the PSE, as well as a number of components inside the PD. Those components do not have perfectly matched resistance, nor does the cable for that matter. The specific combination of PSE, cable and connectors, and PD, causes the total current to divide unevenly between the two pairs of the same polarity. This is referred to as pair-to-pair current unbalance.

Figure 22 shows a schematic overview of pair-to-pair current unbalance. The currents in the two positive pairs are named i_1 and i_2, the currents in the negative pairs are named i_3 and i_4. In a 4-pair system, there is independent current unbalance between the currents in the positive pairs and the currents in the negative pairs. So it is entirely possible that in a running system, four different currents are flowing through the pairs. At any given time, the total positive currents and negative currents are of course equal: $i_1 + i_2 = i_3 + i_4$.

For example, if a system has a total current of 1.2 A, it may be that one positive pair carries $i_1 = 0.55$ A, and the other positive pair carries $i_2 = 0.65$ A, due to current unbalance. The negative pairs may carry $i_3 = 0.7$ A and $i_4 = 0.5$ A in that same system.

In the 802.3bt project, a large amount of effort was spent on establishing reasonable current levels for unbalance based on research into component tolerances and cable properties. The project also had to establish a way to express current unbalance limits using parameters that are observable at the PI. This was difficult because most of the unbalance is caused by components inside of the PSE and PD — something the 802.3 standard doesn't specify.

Based on the developed worst-case components and system model, the highest permitted current unbalance was established. This parameter is called $I_{Unbalance-2P}$ for the PSE, and $I_{Unbalance\_PD-2P}$
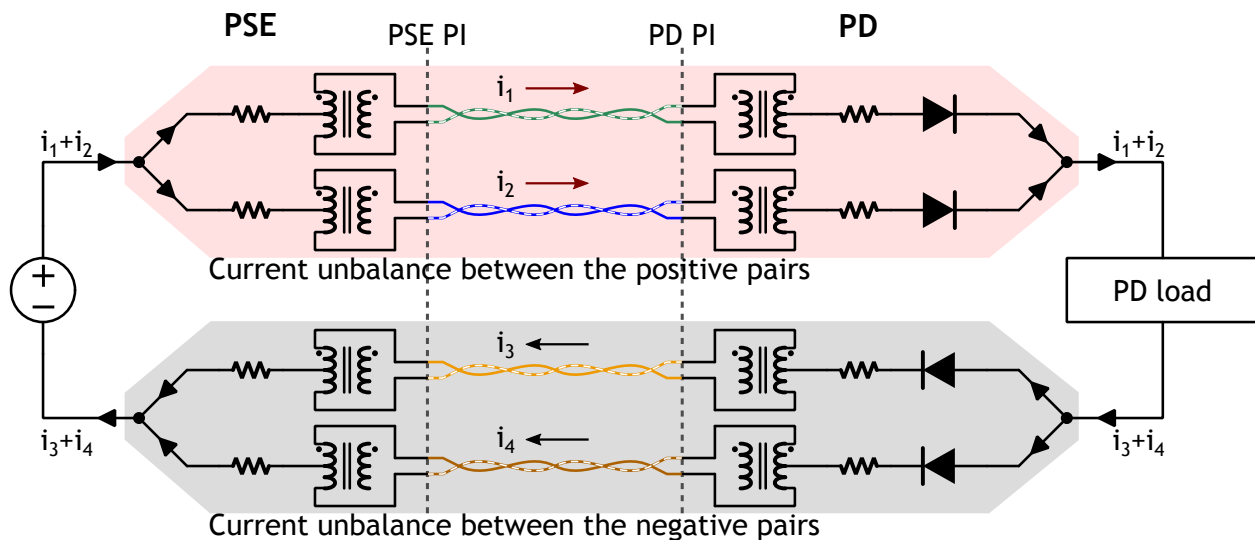
Figure 22: Overview of pair-to-pair current unbalance in 4-pair systems

for the PD. The value depends on the assigned Class. The value of $I_{Unbalance-2P}$ and $I_{Unbalance\_PD-2P}$ is the same (for each Class) and are shown in Table 15. These numbers are the highest amount of current that may flow in any given pair.

If the PSE passes the PSE unbalance requirement, the PD passes the PD unbalance requirement, and a compliant cable is used, the resulting system has pair currents below those shown in Table 15.

9.1. PSE current unbalance

There are two distinct requirements on the PSE relating to current unbalance. First, a PSE is required to support a system with an unbalance current 10 mA higher than what is defined in Table 15. This creates a small amount of margin. The parameter that governs this is $I_{Con-2P-unb}$.

Second, the standard also defines a test to verify that the PSE's contribution to unbalance does not exceed the limits listed in Table 15. This is done by defining a verification circuit, consisting of resistors, that represent the worst-case unbalanced link section and PD. By measuring the cur-

Table 15: Pair current limits for PSEs and PDs due to pair-to-pair current unbalance

Assigned Class	$I_{Unbalance-2P}$ and $I_{Unbalance\_PD-2P}$ (mA)	Highest pair current with perfect balance (mA)
1 to 4	no unbalance requirement	—
5	545	450
6	677	600
7	779	721
8	933	865

rent through all four pairs of the PSE, the designer can verify if the PSE's unbalance contribution is within the limits.

The PSE unbalance verification circuit is defined at two points: a "low link section resistance" and a "high link section resistance". The test resistance values depend on the assigned Class and the PSE should be tested for all the Classes it supports and for the low and high link section resistance for each.

9.2. PD current unbalance

The PD current unbalance requirements are specified similarly to the PSE's contribution to unbalance requirement. A verification circuit is defined, consisting of resistors, that represents the worst-case unbalanced link section and worst-case unbalanced PSE. By connecting the PD to this verification circuit and measuring that the current in any pair is below $I_{Unbalance\_PD-2P}$, the design can be checked for compliance with the unbalance requirements.

Unlike the PSE's verification circuit, the PD's unbalance verification circuit is defined over a continuous range of resistances. PDs that use diodes for rectification generally have the highest amount of current unbalance, resulting from the non-linear way diodes behave when conducting current in parallel. Because of this, PD designers should not assume that the worst-case unbalance occurs at the highest source resistance, but rather test for a suitable number of points in the source resistance range.

Thermal effects can have significant influence on the PD components that cause unbalance and may cause current unbalance to increase or decrease over time. Potential internal or external thermal influence should be taken into account during the design and current unbalance testing.

10. Maintain Power Signature (MPS)

> PSE: 145.2.10 PSE Maintain Power Signature (MPS) requirements
> PD: 145.3.9 PD Maintain Power Signature

When a PSE is supplying power to a PD, the PSE must monitor the current draw in order to make sure that the PD is still connected. The minimum current that the PD must draw to avoid being disconnected is named the Maintain Power Signature (MPS). The PSE is required to remove power when the MPS is absent for at least 400 ms, ensuring that disconnected cables do not remain powered.

The minimum current a PD must draw, called I_{Port_MPS} in the standard, is 10 mA for PDs assigned to Class 1 through 4, or 16 mA if the PD is assigned to Class 5 through 8. Many PDs will have a continuous current draw higher than this. Class 1 through 4 PDs that draw more than 0.6 W continuously and Class 5 through 8 PDs that draw more than 1 W continuously meet the MPS requirements. Those power levels are calculated assuming a PD PI voltage of 57 V, which is the worst-case voltage when it comes to meeting the MPS requirements for a constant-power PD.

Some PDs have low power modes, or sleep modes, with a power level that does not meet the MPS requirements. Such PDs would need to intentionally draw a higher current in order to remain

Table 16: PD MPS requirements and power consumption

PD Type	Assigned Class	PSE Type	Minimum current I_{Port_MPS}	Minimum on time T_{MPS_PD}	Maximum off time T_{MPDO_PD}
1, 2	all	all	10 mA	75 ms	250 ms
3, 4	1 to 4	1, 2	10 mA	75 ms	250 ms
	1 to 4	3, 4	10 mA	7 ms	310 ms
	5 to 8	3, 4	16 mA	7 ms	310 ms

powered, defeating the whole point of having low power standby mode. This was addressed in the 802.3af standard by allowing PDs to draw a pulsed MPS current (see Figure 23). A PD would still meet MPS requirements if the current draw was more than I_{Port_MPS} for at least T_{MPS_PD}, followed by a period of no longer than T_{MPDO_PD}, during which the current consumption does not need to meet I_{Hold}. The values for these parameters are listed in Table 16.

The minimum MPS consists of a pulsed current where the amplitude is dependent on the assigned Class of the PD and the on time and off time are dependent upon the Type of the PSE to which the PD is connected. Type 1 and Type 2 PDs are required to produce a minimum pulse of 10 mA for at least 75 ms with no more than 250 ms between the pulses. This translates to a power consumption of 125 mW from the PSE assuming the operating voltage is 54 V. A PD connected to a Type 1 or Type 2 PSE is also required to support *another* form of MPS measurements named "AC MPS". This involves presenting an impedance of no more than 26.3 kΩ continuously, consuming another 110 mW.

In order to further reduce minimum standby power consumption for PoE systems, Type 3 and Type 4 PDs can make use of optimized MPS timings, when connected to a Type 3 or Type 4 PSE. PDs assigned to Class 1 through 4 must draw a current of 10 mA for at least 7 ms with no more than 310 ms between pulses. This translates to an average power consumption of 12 mW, about 1/20[th] of the Type 1 / Type 2 minimum average power consumption. For PDs assigned to Class 5 through 8, the amplitude of the pulse increases to 16 mA while the timing remains the same, resulting in a minimum average power consumption of 20 mW (PD input power). These requirements are summarized in Table 16.

There are a few important points to take away from this summary. First, Type 3 and Type 4 PDs must still support both the 'long' timings, and the AC MPS requirements, when they are connected to a Type 1 or Type 2 PSE. The PD can recognize the PSE Type by the length of the first class event (see Section 5.3). Second, PDs that request a high power Class (5 through 8) may reduce the amplitude of their MPS when assigned to Class 1 through 4. This allows them to further reduce their standby power consumption at the cost of design complication. Note also that if PDs change their negotiated power level through DLL, the assigned Class may change, which can have an effect on the minimum required current level I_{Port_MPS}.

The PSE current level and timing have been set to take measurement accuracy into account. While Type 1 and Type 2 PSEs were required to keep the PD powered if the measured on time was greater than 60 ms, this time has been reduced to 6 ms for Type 3 and Type 4 PSEs. Note that this shorter time is inherently backwards compatible with Type 1 and Type 2 PDs that produce a 75 ms on time. Type 3 and Type 4 PSEs must remove power when the amplitude of the port

Figure 23: PD MPS requirements

The MPS requirements determine the minimum amount of power the PD needs to draw in order to remain powered. A PD connected to a Type 1 / Type 2 PSE uses the 'long' timings, but also needs to satisfy the AC MPS requirement, which requires it show show DC resistance below 26.3 kΩ. With a PD voltage of 54 V these requirements result in a minimum power consumption of 210 mW.

A Type 3 / Type 4 PD connected to a Type 3 / Type 4 PSE however may make use of the 'short' MPS timings, and is not required to show the AC MPS DC resistance. When assigned to Class 1 to 4, the minimum power consumption (again, assuming 54 V at the PD) is 12 mW. When the assigned Class is 5 to 8, a higher current needs to be drawn, resulting in a minimum consumption of 20 mW.

Table 17: PSE I_{Hold} and $I_{Hold-2P}$ MPS current levels

Assigned Class, power mode	$I_{Hold-2P}$			I_{Hold}		
	min	max		min	max	
Class 1 to 4, 2-pair power	4	9	mA	4	9	mA
Class 1 to 4, 4-pair power	2	5	mA	4	9	mA
Class 5 to 8, 4-pair power	2	7	mA	4	14	mA

current is below 4 mA for longer than 400 ms. The PSE may remove power when MPS is absent for more than 300 ms.

10.1. PSE MPS measurements

At very low current levels (such as MPS), there is no feasible way for a PD to guarantee any form of pair-to-pair current balance. This means that all the current could flow on Alternative A, or all the current could flow on Alternative B, or that the current is evenly divided between the two. The pair-to-pair current unbalance can also dynamically change while the PD is on due any number of reasons (e.g. thermal drift). The PSE MPS requirements are designed such that, for 4-pair operation, there are two possible measurement strategies for the PSE to determine if the current being sourced meets the MPS requirements. Figure 24 shows an overview of the two methods.

In case of 2-pair powering, the PSE compares the current of the active pairset, $I_{Port-2P-pri}$, against $I_{Hold-2P}$. The limits for $I_{Hold-2P}$ are given in Table 17.

In case of 4-pair powering the PSE can measure the total current (sum of both pairsets: I_{Port}) and compare this against I_{Hold}, see Figure 24. Alternatively, the PSE can measure the current of each pairset individually. These currents are called $I_{Port-2P-pri}$ and $I_{Port-2P-sec}$. The greater of these two can be compared against $I_{Hold-2P}$.

The PSE is required to recognize the MPS as **present** when the corresponding PI currents exceed both $I_{Hold-2P}$ max and I_{Hold} max. MPS is to be recognized as **absent** when the corresponding PI currents are below $I_{Hold-2P}$ min and I_{Hold} min. In all other conditions the PSE may determine the MPS as either **present or absent**. The PSE is not required to use both methods, the requirements are such that the methods never produce a conflicting answer.

If there is sufficient current present for at least T_{MPS}, followed by no more than T_{MPDO} where current may be absent, power is maintained. If the PD fails to meet MPS for at least T_{MPDO}, the PSE is required to remove power.

10.2. PD MPS design consideration

There is an important design consideration for PDs that want to make use of the optimized Type 3 / Type 4 MPS timings. The Type 1 / Type 2 timings have 15 ms of margin between the PD minimum on time ($T_{MPS\_PD} = 75$ ms), and the PSE MPS valid time ($T_{MPS} = 60$ ms). With the short
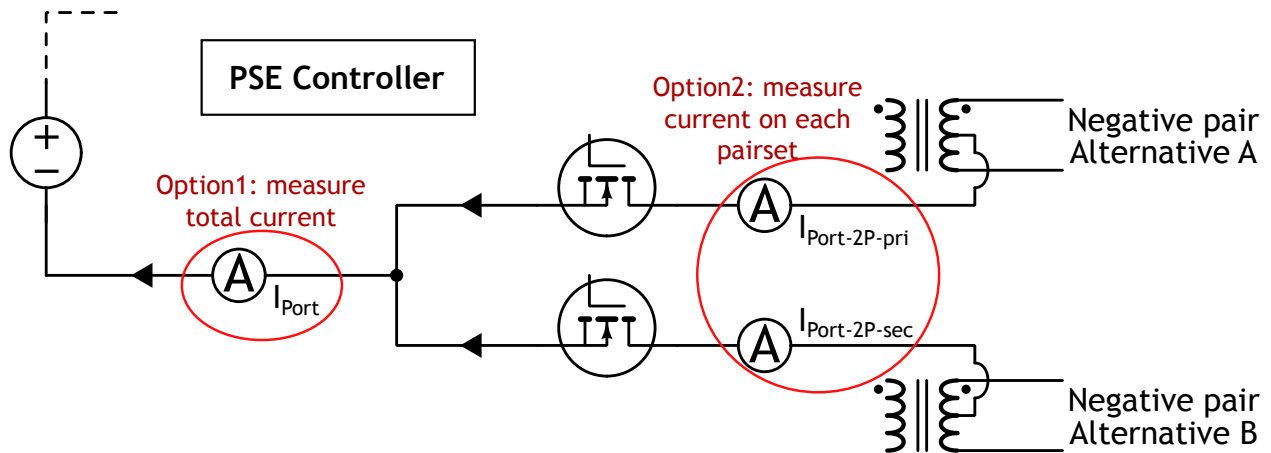
Figure 24: PSE MPS presence / absence determination methods

timings, that margin is only 1 ms as the PD minimum on time is 7 ms and the PSE MPS valid time is 6 ms. A PD with very low power consumption that intentionally draws MPS pulses to remain powered, usually has a current sink or resistor that it can control. When that circuit is activated at low load, the resulting current is initially mostly sourced from the bulk capacitor (indicated by the grey current flow in Figure 25) and not from the PSE, because the bulk capacitor current path is lower resistance than the path from the PSE. As the voltage of the bulk capacitor drops, more current is drawn from the PSE until finally the voltages have equalized and all current is drawn from the PSE.

This causes current waveforms to appear at both PSE and PD PI that look like those in Figure 26. As can be seen in the example, even though the MPS current generating circuit may be activated for the correct amount of time, the time that the current sourced from the PSE effectively meets the MPS requirements may very well be lower. With the 'short MPS' timings, this becomes an area of attention.

PD designers must test their MPS circuits with a PSE connected through a high resistance channel. A worst-case channel ($R_{Chan} = R_{Ch}$) will exhibit the current 'stealing' by the bulk capacitor effect most strongly. PDs can employ a number of methods to meet T_{MPS}, as measured on the PSE PI, such as drawing a higher MPS current, keeping the MPS circuit on for a longer amount of time, or by forcing the current to be sunk from the PSE, rather than the bulk capacitor.
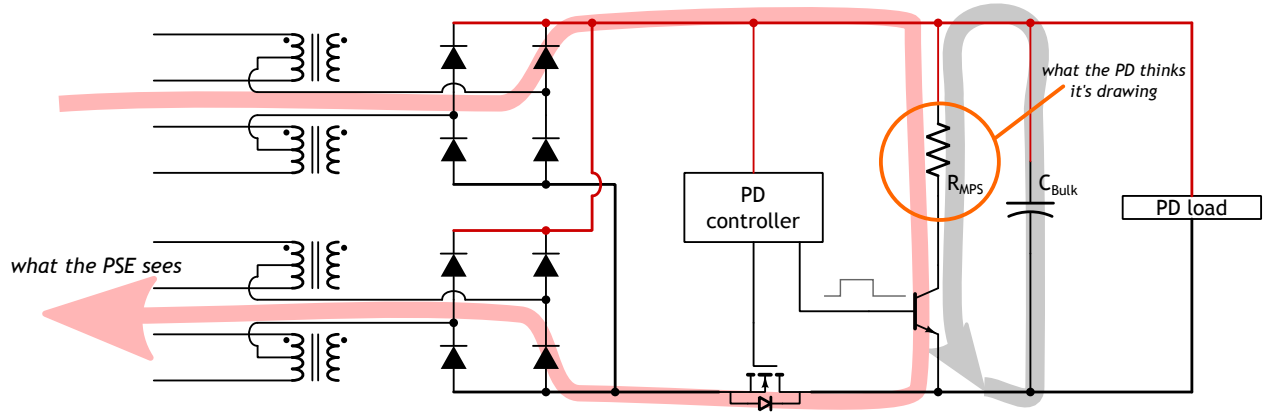
Figure 25: Example PD MPS circuit that partly draws current from the bulk capacitor
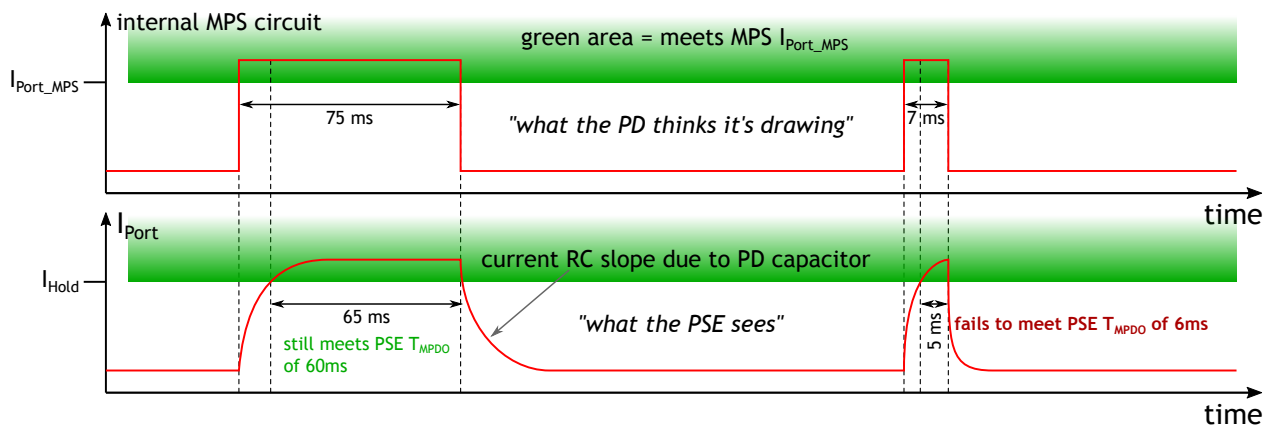

Figure 26: PD capacitor causing MPS pulses to be narrower than expected due to RC effect

11. LLDP Power over Ethernet TLV

> Clause 79 in 802.3 and in 802.3bt

The LLDP protocol serves to manage individual Ethernet links. The protocol is based on TLVs (Type/Length/Value). One of these TLVs is reserved for "organizationally specific" purposes, one of which is assigned for use by IEEE 802.3. Within this TLV, subtypes are defined that each control a set of link properties. One of these subtypes is used for Power over Ethernet management and another is used for Power over Ethernet measurements.

Within the PoE subtype, Clause 79 defines the fields in Table 18. A number of these fields are used to support power negotiation between a PSE and a PD, named DLL (Data Link Layer) classification.

Table 18: Description of LLDP fields

Field name	Size	Description
MDI power support	1 byte	This field indicates if the device is a PD or PSE, if it supports Power over Ethernet, and whether the PSE has the ability to control which pairs provide power.
PSE power pair	1 byte	This field indicates which of the pairsets is providing power. If both pairsets provide power, either pairset may be indicated.
Power class	1 byte	This field indicates the requested Class of the PD as detected by the PSE. When sent by a PD, the value is undefined.
Type/source/priority	1 byte	This field indicates the Type of the device (Type 1 or Type 2) and if it is a PD or a PSE. It also indicates details of where the PD is receiving power from, and what the power source of the PSE is, as well as an indication for the powering priority as indicated by the PD. A new bit was added by 802.3bt to this field, to indicate if the PD supports 4-pair power.
PD requested power value	2 bytes	This field is part of DLL classification and is used by the PD to request a certain amount of power, expressed in $1/10^{th}$ of a Watt. When sent by a PSE, this field contains the last value the PSE received from the PD. That allows the PD to verify that the mechanism is synchronized. PSEs and PDs of all Types make use of this field.
PSE allocated power value	2 bytes	This field is part of DLL classification and is used by the PSE to grant a certain amount of power, expressed in $1/10^{th}$ of a Watt. When sent by a PD, this field contains the last value that the PD received from the PSE. This allows the PSE to verify that the mechanism is synchronized. PSEs and PDs of all Types make use of this field.

Field name	Size	Description
PD requested power value Mode A	2 bytes	This field is used exclusively by dual-signature PDs to negotiate power on a per-pairset basis. Single-signature PDs set this field to zero.
PD requested power value Mode B	2 bytes	This field is used exclusively by dual-signature PDs to negotiate power on a per-pairset basis. Single-signature PDs set this field to zero.
PSE allocated power value Mode A	2 bytes	This field is used only by PSEs that are powering a dual-signature PD in 4-pair mode. In all other cases the field is set to zero.
PSE allocated power value Mode B	2 bytes	This field is used only by PSEs that are powering a dual-signature PD in 4-pair mode. In all other cases the field is set to zero.
Power status	2 bytes	This field provides classification status information for Type 3 and Type 4 devices. It contains if the PD is single- or dual-signature, and if it is 2-pair or 4-pair powered, which pairs power is being provided on (similar to PSE power pair), and the requested Class of the PD (when sent by a PD), and the assigned Class of the PD (when sent by a PSE).
System setup	1 bytes	This field indicates static properties of Type 3 and Type 4 devices, it communicates the Type and signature configuration, as well as indicates if a dual-signature PD's load is electrically isolated.
PSE maximum available power	2 bytes	This field indicates the highest power allocation the PSE would be able to grant, at the moment of sending. A PD can read this field to find out what power request (through PD requested power value) may succeed, rather than guessing.
Autoclass	1 bytes	This field contains a number of bits that allow the PD to request Autoclass, and the PSE to confirm that it has completed the measurement.
Power down	3 bytes	This field allows a PD to request to be powered down. It also allows the PD to set a wakeup timer in the PSE, after which the PSE can re-apply power.

11.1. Power via MDI Measurements TLV

Next to the "Power via MDI" TLV, an optional TLV, "Power via MDI measurements" is also defined. With this TLV the PSE and PD can exchange information about electrical conditions at their respective PI. Four different measurements are supported: voltage, current, power, and energy. The devices independently can indicate which measurements are supported, which measurement they request from the link partner, what the quality of the measurement is (expressed in an uncertainty figure), and finally the measurement results itself.

The final two bytes of this TLV contain the PSE power price index. This integer number allows the PSE to communicate the current relative price of electricity compared to what the PSE considers a nominal price point.