

April 12, 2006

Simple circuit design tutorial for PoE applications

By Robert Mayell, Staff Applications Engineer, Power Integrations

As the number of Power Over Ethernet (PoE) powered device (PD) applications grows, market pressures are driving designers to lower the cost and complexity of the DC-DC converters that power them, while improving their performance. This paper provides simple circuit designs that meet those goals. It presents a discrete PD interface circuit that has been shown to work in [University of New Hampshire Interoperability Consortium \(UNH-IOC\)](#) tests. That circuit can be implemented with only a few low-cost components. It gives practical design hints and tips on designing a PD DC-DC converter stage. It also presents a galvanically isolated, Flyback DC-DC converter example.

Lastly, it addresses designing PDs to operate from alternate power sources, such as AC power adapters. Since PoE enabled PDs will become consumer commodity items, the simplicity, cost-effectiveness and performance of the circuits presented in this paper make them attractive for most PD applications.

The current specification and what it mandates

The specification that describes the delivery of power over CAT-5 Ethernet cabling is an amendment to part three of the main IEEE Ethernet specification. Specifically it is: [IEEE 802.3af Part 3: Carrier Sense Multiple Access with Collision Detection \(CSMA/CD\) Access Method and Physical Layer Specifications](#) Amendment: Data Terminal Equipment (DTE) Power via Media Dependent Interface (MDI). The first piece of information provided by the spec, that is relevant to PD circuit design, describes how the PSE can provide the power to the Ethernet cabling system (see Figure 1). The system on the left hand side of Figure 1 uses a PoE enabled switch/router to provide power--through the patch panel--to the PDs.

The system on the right hand side uses an older switch that does not provide power, so the power must be injected somewhere between the switch and the patch panel. Since the power is added at the starting "end" of the enabled switch, it is referred to as an "endpoint" system, versus a midspan, which injects the power somewhere in between the switch and the patch panel.




Figure 1. PSE power injection schemes: Endpoint (left) vs. Midspan (right)

Essentially, midspan PSEs are merely a way to allow PoE to be added to older systems without replacing the switch hubs (with newer, PoE enabled units). The current specification only allows power on two of the four wire pairs in the CAT-5 cable. As a general practice, Endpoint PSEs place their power onto the data pair of wires in the CAT-5 cable (see Figure 2), while midspan PSEs are restricted to using the spare pair of CAT-5 wires (see Figure 3 and 802.3af, pages 29 and 30). The specification also allows three different power connection options (802.3af, page 31), which determine the number and orientation of diodes found on a PD front-end.




Figure 2. Endpoint PoE configuration with PSE injecting power onto data pair

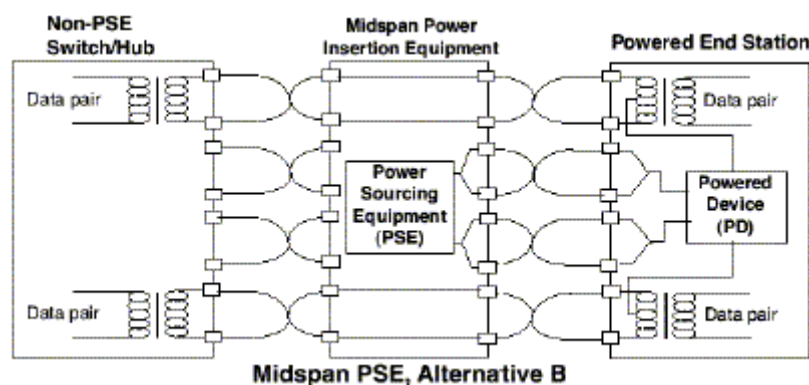
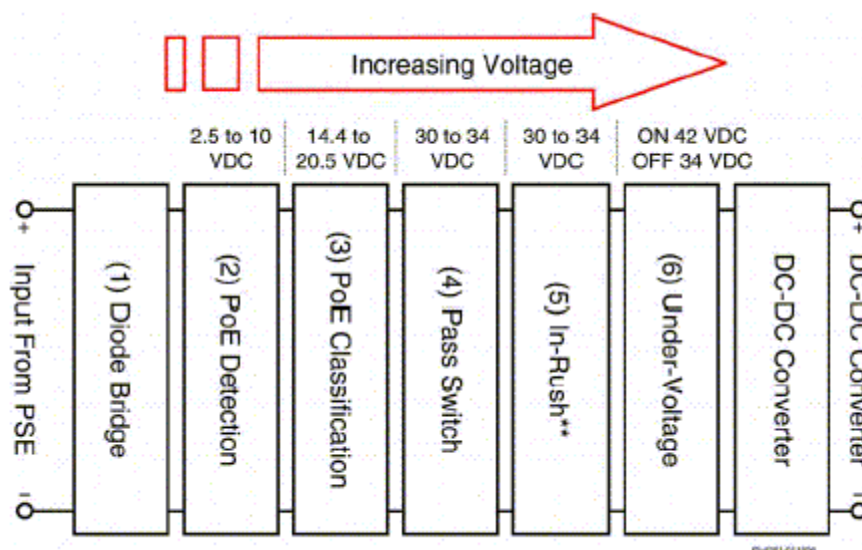


Figure 3. Midspan PoE configuration with PSE injecting power onto spare pair

The front-end interface of powered devices

How powered devices must interface with power sourcing equipment

When a PD is connected to a powered Ethernet cable (link), a series of steps must successfully occur as it interfaces with the PSE, in order for the PSE to provide power to the PD. Figure 4 is a combination block and sequencing diagram that shows the steps a PD goes through when interfacing with a PSE.



****Inrush not required**

Figure 4. Block/sequencing diagram of a PD successfully interfacing with a PSE

The input diode bridge (see Figure 9) protects the circuitry within the PD from being connected to a reverse polarity input voltage. The first phase (detection) of the powering sequence occurs when the PSE polls the newly connected device to see if it provides the correct impedance signature. The PSE accomplishes this by ramping up a current limited (5 mA) detection voltage (from 2.5 V to 10 V) across the designated pairs of CAT-5 wires (at about a 2 ms repetition rate) and measuring the voltage and the current at the end of the ramp time. If the PSE detects the proper signature impedance (802.3af, pages 38 and 39, Table 33-2), it determines that there is a valid PD at the end of the link. The PSE then proceeds to the next step in the process -- Classification.

In the Classification step, the PSE applies more voltage (a 0 to 20.5 V ramp) to the link, and measures the current drawn, in order to determine the classification (its power requirement; 802.3af, page 53, Table 33-10) of the PD. Once the PSE has classified the PD, it is supplied with full operating voltage so that it can connect its pass switch and start up its DC-DC converter.

In-rush protection was required in some older, pre-PoE systems. However, in the newer systems, in-rush current is so low that there is no need to have any current limiting components in the PD.

Under-voltage lockout (UVLO) is required due to the length of CAT-5 cable (100 M, or about 330 ft) over which power must be supplied. The voltage drop across a 100 M cable can be up to 8 V when the maximum available power (12.95 W) is drawn by a PD. The DC-DC converter's UVLO circuit must ensure that the voltage at the PD is high enough (≥ 42 V) that it will not drop below the minimum working voltage (34 V) once load current is drawn from the cable. This prevents the DC-DC converter within the PD from starting up and then shutting down repeatedly as power is initially applied to the link.

From detection to drawing power: getting past the handshake

There are a number of IC solutions available that will interface with a PSE to deliver an acceptable signature impedance, draw the correct classification current, and drive a pass switch device. However, all of those functions can easily be done with less than 17 discrete components.

Detection: Who approaches? Identify yourself!

First, the detection impedance requires only one resistor with a value between 19 k Ω . and 26.5 k Ω .. Resistor R51 (Figure 5) draws a current close to the minimum expected of a valid impedance, to minimize consumption after the DC-DC converter is working.




Figure 5. The detection impedance resistor (R51) of a simple, Class 0, PD interface circuit

Classification: How much power will you be needing? During the classification phase, the PSE uses the current drawn by the PD to determine how much power it will require. For Class 0 (the general, default Class), the detection impedance draws the correct amount of current to classify the PD as a Class 0. This makes the front end of the Class 0 PD the simplest and least expensive solution.

The circuit diagram of the Class 0 circuit (Figure 5) consists of six components that cost about 12 cents (at high volume manufacturing pricing) in materials. A less expensive (but lower efficiency) bipolar transistor based pass switch circuit requires 3 additional parts (to boost the drive into the base of the transistor) but costs only about nine cents in materials.




Figure 6. Detection, Classification and Pass-Switch interface circuit for Classes 1–3

Identifying a PD as a Class 1, 2 or 3 requires a few additional parts, as can be seen in Figure 6. The 10 components (VR1, Q1, Q2, Q3, Q4, R2, R3, R4, R5 and U1) that are located between the signature impedance resistor and the MOSFET pass-switch connection circuit only conduct current once the PSE's probing voltage exceeds the Zener voltage (11 V) of VR1. The PSE ramps the voltage across the link from 0 to 20.5 V. It measures the current drawn by the PD while the voltage is ramping from 14.5 to 20.5 V. During this phase of the ramp, the PSE expects to measure a current within the range of 9 and 30 mA to detect a valid PD of Class 1, 2 or 3 (see the table in Figure 6).

The classification circuit works as follows: Transistors Q1, Q2, R2 and R3 form a 350 μ A current source (set by the value of R3) that biases U1. Transistor Q3 and U1 form the current source that sets the classification current, based on the value of R4. By changing the value of R4, the circuit will conduct the right amount of classification current for a Class 1, 2 or 3 device. Once classification has occurred and the pass-switch is connected, R5 biases Q4 on, which shunts the 350 μ A current source to ground. This disables U1 and Q3 to minimize the power consumption of the front-end once the DC-DC converter begins operating.

Pass switch connection: power has been applied

Once the PD has been properly classified, the PSE is then ready to apply full power to the link. As the PSE ramps the power up to full voltage, VR2 begins conducting at about 28 V, which then turns on Q5 and Q4. Resistor R7 ensures that the current that flows through VR2 develops sufficient voltage across the gate to source junction of Q5 to fully turn it on. Zener diode VR3 ensures that the maximum gate voltage of Q5 isn't exceeded when the supplied input voltage is at its highest (57 V). Resistor R6 limits the current through VR2 to minimize its dissipation and power loss, and to protect it from over-heating. The cost of the 16-component Class 1–3 discrete front-end is about 20 cents (in high volume).

The cost of gaining access: the path of least expenses

As market pressures make cost the major factor in choosing components for PD circuitry, the discrete interface circuit described here will be the benchmark, since lower cost solutions do not exist.

Designing the PD's DC-DC converter

Converter Topology Choices: Flyback, Forward or Buck

Depending on the voltage and current requirements of the PD, there are three converter topologies that could be used in a powered device DC-DC converter: flyback, forward or buck.

Description	Comments
Flyback	Lowest cost for output currents < 6 A
When to Use	Typically used for output voltages (> 2.5 V) Recommended for output voltages (> 12 V) For applications requiring lowest cost
Advantages	No output inductor required. Output isolated from input.
Disadvantages	Higher output ripple current (higher output capacitor cost) Medium efficiency (due to higher peak and rms currents)
Forward	Lowest cost for output currents > 6 A
When to Use	Typically used for output voltages (< 12 V) Recommended for output voltages (< 2.5 V) For

	applications requiring highest efficiency
Advantages	Low output ripple current (lower output capacitor cost) Higher efficiency (due to lower peak and rms currents) Can utilize synchronous rectification. Isolated output.
Disadvantages	Requires output energy storage inductor (higher cost)
Buck	Usable if isolation not required and output current < 3 A
When to Use	Low power, one output, and low efficiency is acceptable
Advantages	Simple, low parts count. Uses off-the-shelf inductors
Disadvantages	Variable switching frequency required for low output voltages (< 12 V). One output only (multiple outputs not recommended). Switch carries full output current.

Table 1: Flyback, Forward and Buck: Benefits and Disadvantages

The benefits and disadvantages of each topology

Since cost is such a significant factor in PD design, the least expensive DC-DC converter will typically be a Flyback, unless galvanic isolation is not required (a Buck or non-isolated Flyback could be used), or the load current approaches 6 A (a Forward converter would be preferred).

Since most PD applications require input voltages and currents that can efficiently and cost effectively be delivered by Flyback converters, this section will focus on the operation and design of Flyback converters.




Figure 7. Basic Flyback Energy Delivery (Discontinuous Conduction Mode shown)

For a single switching cycle of a Flyback converter's normal operation: the controller turns the MOSFET switch on, creating a path for current (I_D) to flow through the primary winding of the transformer and the MOSFET. Due to the phasing of the transformer primary and secondary windings and the output diode orientation, no current flows in the secondary while the switch is on, and energy is stored in the core of the transformer. The current increases linearly until the controller turns the switch off. As the switch turns off, the voltage across the transformer windings reverse, which forward biases the output diode. The energy that was stored in the transformer core is then transferred to the output capacitor (which is connected across the secondary winding) as current (I_S) flows through the output diode.

Flyback design options: RCD and Zener diode clamps

Every time the controller turns the switch off, the leakage inductance of the transformer generates a voltage spike at the junction between the transformer and the switch, as the voltage inversion occurs. A clamp network is used to protect the switch from the potentially destructive energy in the voltage spike. In Flyback converters, one of two drain-node clamp circuits are typically used to protect the MOSFET from the Flyback voltage spike: a Resistor, Capacitor and Diode (RCD), or a Zener diode.

The voltage rating of the Zener diode is chosen so that it does not conduct during normal operation, but only when a line surge or load transient induces a larger-than-normal spike. With a Zener clamp, the converter usually operates more efficiently. An RCD network is used when EMI noise generation is a concern, since it clamps the Flyback voltage spike each time the switch turns off. The diode in the RCD clamp should have about a 500 ns recovery time, so that some of the energy stored in the leakage inductance can be recycled as it recovers. A series resistor limits the reverse current through the diode. The values of the clamp resistor and capacitor are typically obtained through an iterative process of taking measurements and adjusting values until the spike is sufficiently clamped while the temperature rise of both components is kept to a minimum. This is required because the values of resistance and capacitance will primarily be determined by the parasitic capacitances and leakage inductance of the transformer, which cannot be easily predicted before a number of transformer samples have been built, tested and characterized.

Flyback transformer core selection guidelines The following table lists the transformer cores most commonly used in low-power, DC-DC Flyback converters.

Description	Comments
EFD EFD10, EFD12 EFD15, EFD20	Low profile. Rectangular center leg, hence higher mean length per turn (MLT). Terminating multi-filar windings to EFD bobbin pins is difficult.
ER ER9.5, ER11, ER14.5	Low profile. Round center leg, hence lower MLT. Large center leg gives gives high inductance with few turns, hence lower leakage inductance. High pin count bobbins make winding termination easy.
RM RM6, RM10	Large center leg. Terminating multi-filar windings to RM bobbin pins is difficult. RM bobbins have little room for meeting safety spacing requirements. Cores are expensive.
PR PR14x8	Large center leg. Terminating multi-filar windings to PR bobbin pins is difficult. PR bobbins have little room for meeting safety spacing requirements. Cores are expensive.
Planar ELP18/10/4, E18, E22, ER18, ER23	E series has a square center leg, hence high MLT. ER series has a round center leg, hence lower MLT. Not cost effective at PoE power levels.

Table 2: Flyback Transformer Core Selection Guidelines

Flyback transformer construction guidelines The following table lists the most important aspects of low-power, DC-DC, Flyback transformer design.

Description	Comments
Core Flux	Maximum Flux Density (B_M), measured in Gauss or Tesla
BM	(higher efficiency) $1500 < B_M < 2500$ (smaller core size)
Core Gap	Gap length (L_G), typically measured in millimeters
LG	$L_G > 0.1$ mm (the minimum, reliable gap size)
Windings	How many turns and layers (of turns) per winding
Number of turns	Few turns limits leakage inductance and conduction losses
Split windings	Reduces leakage inductance
Multi-filar ≤ 4	Limited by the number of termination pins. Reduces Skin Effect
Copper Foil	Use only if > 4 strands of multi-filar wire are required

Table 3: Flyback Transformer Construction Guidelines

Flyback output diode selection guidelines The following table lists the most commonly used output rectifier diodes for low-power, DC-DC, Flyback converter.

Description	Comments
Low-Drop Schottky	For low output voltage and high efficiency applications
SL1x, SL2x, SL4x	Low forward voltage drop: $0.44 V_{FORWARD}$, rated for 1 A, 2 A and 4 A, and 20, 30 and 40 V of reverse voltage
Schottky	Reliable Schottky diodes with $V_{REVERSE}$ up to 100 V
SS1x, SS2x	Forward voltage drop range of 0.5 V to 0.75 V
Ultra Fast	For high voltage output (> 12 V) applications
ES2x	20 ns recovery time, 50 V to 200 V reverse voltage
BYG22x	35 ns recovery time, 50 V to 200 V reverse voltage

Table 4: Flyback Output Diode Selection Guidelines

Low-drop Schottky diodes are typically used for output voltages ≤ 3.3 V. Standard Schottky diodes work well for an output voltage range of 5 V to 9 V. Ultra Fast diodes are typically used when the output voltage is ≥ 12 V.

Flyback output capacitor(s) selection guidelines

The following table lists the most commonly used output filter capacitor types in low-power, DC-DC, Flyback converters.

The capacitance value will depend on the peak secondary current, the part's ripple current rating, and how many capacitors are used in parallel.

Description	Comments
Tantalum	A common choice for DC-DC Flyback converters. Low ESR and high capacitance values. Require significant voltage derating (30–50 %). Short circuit failure mode.
T491 & T495 series, TE-L series	Kemet and Panasonic, respectively
Aluminum Organic	Super low ESR. Expensive, but coming down. Require no voltage derating. Open circuit failure mode.
A700 series, 6PTB series	Kemet and United Chemicon, respectively
Ceramic	Very low ESR, but usually low capacitance also. Low ESR can cause loop stability problems.
SMD Electrolytic	A good fit for most applications, but expensive. Not as reliable as AO or ceramic. No loop stability problems.

Table 5: Flyback Output Capacitor Selection Guidelines

Tantalum capacitors are most commonly used. SMD Electrolytics may provide a good cost versus reliability compromise--if space allows--until the cost of Aluminum Organic capacitors comes down further. Ceramics can be used, but loop compensation must be given careful attention to ensure stable operation.

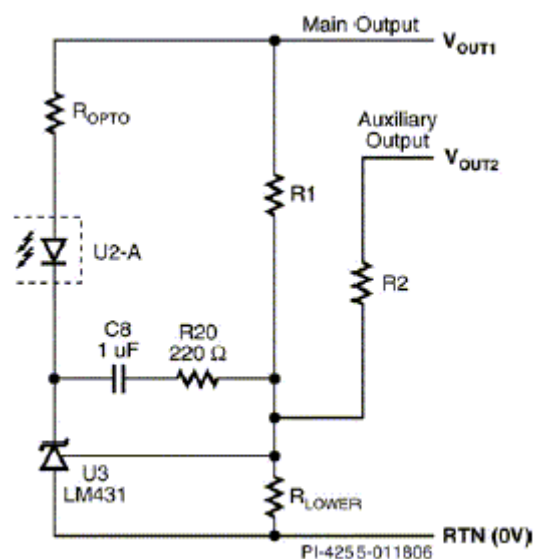


Figure 8. A typical TL431 circuit's resistor divider & frequency compensation components

Feedback design: TL431 resistor divider network

A resistor divider network is used to deliver a portion of the output voltage to the reference (control) pin of the secondary reference IC. The resistance values used are calculated so that the current that flows through the lower resistor (R_{LOWER} in Figure 8) is about 100 μA . Equations are shown in the session presentation.

Feedback design: TL431 frequency compensation

The purpose of the reference IC is to provide high gain for DC and low frequency feedback signals. Capacitor C13 (see Figure 9) rolls off the gain of U3, as frequency increases.

Flyback quick design checklist

Three critical design parameters must be verified in a Flyback converter: thermals, maximum drain voltage, and peak drain current.

First, it is imperative that no power handling components get so hot that they fail, or cause other components or circuits to operate incorrectly, even during abnormal conditions. Second, the designer needs to verify that the peak drain voltage does not exceed the maximum voltage rating of the MOSFET switch. This usually occurs when the converter is

overloaded while operating at the highest input voltage (57 V), just before it loses regulation of the output and goes into self-protection mode. Third, the designer must also verify that the drain current waveform does not show any indication of transformer saturation (decrease in inductance near the peak of the primary current ramp), especially at start-up and high line overloads.

A practical dc-dc converter design example

Figure 9 is the circuit diagram of a typical PD converter. The PoE front-end identifies the PD as a Class 2 device. The DC-DC converter is a Flyback that uses a Zener diode clamp (VR3) on the MOSFET drain-node. It uses a 20 V, 4 A, standard Schottky diode (D1) for its output rectifier, and tantalum output capacitors (C7–C9). The frequency compensated (R11, C13) TL431 (U3) drives feedback through the opto-coupler (U2). Lastly, a Y capacitor and damping resistor (C4 and R7) attenuate common-mode EMI noise.




Figure 9. 6.6 W, single output (3.3 V) Flyback converter circuit diagram

Designing PDs to operate from AC-DC adapters

Keep PD DC-DC design simple: AC-DC operation is temporary

Requests are often made for PDs to be able to operate from a 12 V, AC-DC power adapter in addition to 48 V. This has tempted some PD designers to make the DC-DC converter within their device be able to operate over a wide input voltage range (12–60 V). This makes a DC-DC converter significantly more complex and costly. The most important fact to remember is that operating from an AC-DC adapter is only a temporary, short-term solution that will go away, once all switches, hubs and routers have PSEs built into them. A better solution would be to use a single link, AC-DC, 48 V adaptor, such as the one pictured in Figure 10. This would allow PD designers to keep the DC-DC converter within the PD as simple and low-cost as possible.

The unpowered CAT-5 cable plugs into the input (lower) jack (J1), and the injected 48 V power is placed on the unused (Alternative B) pair of wires that connect to the output (upper) jack (J2). This maintains the current midspan PSE convention, and prevents the 48 V from being connected to the unpowered switch or router. This solution does not include the detection and classification functions specified for a PSE by 802.3af, but that does not matter, since it is a temporary solution that only powers a single port.




Figure 10. The circuit diagram of a single port, midspan (21.7 W) PSE

References:

1. IEEE 802.3af Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Data Terminal Equipment (DTE) Power via Media Dependent Interface (MDI).
2. Application Note AN-31: DPA-Switch DC-DC Forward Converter Design Guide, April, 2003, Power Integrations.
3. Application Note AN-32: TOPSwitch-GX Flyback Design Methodology, December, 2002, Power Integrations.
4. DPA-Switch Family datasheet., July 2005, Power Integrations.
5. Engineering Prototype Report (EPR)–86: 6.6 W, Multi-Class Powered Device (PD) for Power over Ethernet (PoE) Using DPA423G, February 2006, Power Integrations.
6. Design Idea DI-70: DPA-Switch PoE Detection and Classification (Class 0) Interface Circuit, November 2005, Power Integrations.
7. Design Idea DI-88: DPA-Switch PoE Detection and Classification (Class 0-3) Interface Circuit, November 2005, Power Integrations.
8. Design Idea DI-101: DPA-Switch Under-Voltage with Wide Hysteresis, November 2005, Power Integrations.

About the author

Rob Mayell is a Staff Applications Engineer at Power Integrations. He has held a variety of design engineering positions from power supply design to custom IC development in the UK, USA, and France and was an Engineering Manager for Schlumberger prior to joining Power Integrations in 1997. Since this time he has worked on both AC-DC and DC-DC

converters including Power Over Ethernet (PoE) applications. BE Honors Degree in Electrical Engineering from Salford University, Manchester England rmayell@powerint.com