Challenges and solutions for PoE systems in Ethernet switches

By Penny Xu
Applications engineer, PSIL-PI

Introduction

With more devices than ever trying to connect to the Internet of Things, powering these devices has become a hot topic. Power over Ethernet (PoE) is a newer way to provide DC power while also accommodating data through an Ethernet cable. Using PoE power has a lot of benefits:

- Lower bill-of-materials costs because there is no need for a wall adapter or local power supply, nor is there a need for professional installation by an electrician.
- PoE has multiple levels of protection for overload, short circuit and overtemperature. PoE devices perform hand-shaking before delivering power, which prevents damage from incorrect installations or unknown devices.
- With PoE pass-through, it is possible to deliver power anywhere in the network with a star or mesh connection.
- A central controller has access to every PoE port, from which it can obtain system status and diagnose system faults.

The Institute for Electrical and Electronics Engineers (IEEE) has reported that PoE is "growing well beyond shipments of 100-million ports per year." With the ratification of the IEEE 803.3bt standard and the adoption of PoE to be considered for end equipment such as electronic point of sale, professional audio and video, digital displays, and a whole myriad of building-automation solutions, PoE will continue to grow at a fast pace.

Ethernet switches play an important role in PoE systems, as they power powered-device (PD) loads such as Internet Protocol (IP) phones, IP cameras, and wireless access points and speakers. Controlling power-sourcing equipment (PSE) in an Ethernet switch is complicated, however, given the limited capability from the power supply and the multiple power supplies that exist for redundancy in the event of a power outage.

PoE system overview

Figure 1 shows a basic PoE system that consists of a PSE in a switch or hub, a Category-5 cable, and a powered device (PD) in a powered end station such as a business phone. The PSE sends power through the Ethernet cable and the PD receives power from the PSE via the data pairs. Spare cable pairs are also used to transfer power if the end station requires more power than allowed through the data pairs.

Figure 1. PoE System Overview

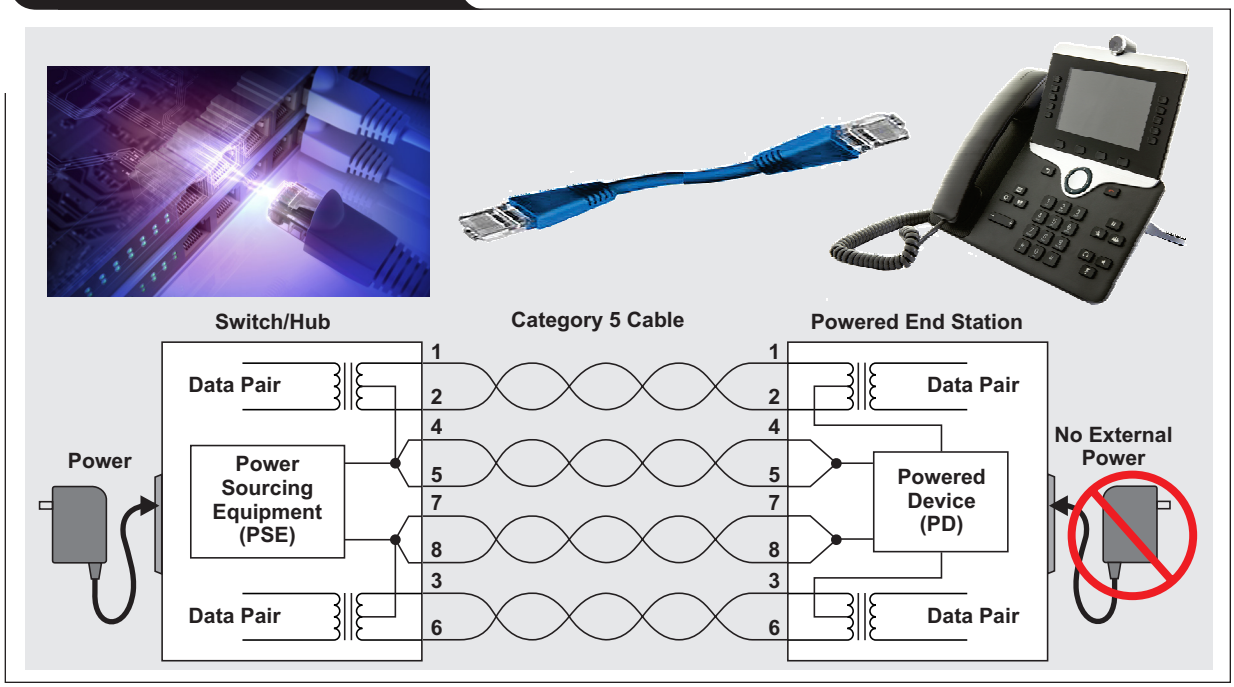

In the latest IEEE 802.3bt standard, power limits have increased to 90 W for sending and 71 W for receiving to PDs across 100 m of a Category-5 cable. The PSE only provides power to valid PDs. The PSE and PD have to finish a negotiation—detection, classification and in-rush—before turning on the port.

Figure 2 shows the mutual identification between PSE and PD. The PSE starts the negotiation process by sending a detection current/voltage signal and calculating the resistance based on the measured voltage/current. The resistance is considered valid if it falls in a range between 19 kΩ and 26.5 kΩ. After successful detection, PSE starts classification by sending voltage signals. The number of class event (1 through 5 events) depends on the PD's class level and PSE-port's power application. At the same time, the PD should present different class resistors on the port to indicate its power level. The PSE measures the current going through the class resistor and recognizes the class level (power level) of the PD.

After a valid classification result (no overcurrent or class mismatch), the PSE turns on the port and starts monitoring the inrush current. If the current remains below the PSE's in-rush current limit (~425 mA) during the Tstart period (the maximum allowed overcurrent time during in-rush), the PSE and PD enter normal operating mode and the PD can pull a full load.

PSEs in Ethernet switches

Figure 3 shows one of the most common PoE systems. The microcontroller (MCU) is dedicated to PSE control through I^2C and communicates with the host's central processing unit through I^2C, serial peripheral interface (SPI), or universal asynchronous receiver transmitter (UART) interfaces. A DC power supply (44-V to 57-V) powers the PSE board and delivers power to PD loads through Ethernet cables.

PSE design challenges

The technical design challenges of the PSE in PoE systems include:

- **Port power management:** The Ethernet switch usually has multiple ports (16, 24 or 48), but the power supply cannot supply full power for all ports. For example, in a 48-port configuration, each load connected to a port can consume as much as 30 W. So the PSE can draw as much as 1,440 W of total power. In order to save costs, most manufacturers do not ship systems with AC power supplies rated to handle a call for full power from all ports. The challenge is how to power as many PDs as possible, limit the power cycling of PDs and keep high-priority ports on in the event of a power shortage.
- **Load-step transitions:** There are different ways to handle load variations with port power management. The most popular method is based on every port's real-time consumed power—in other words, dynamic control. This method is efficient, as there's no wasted power. Getting real-time consumed power is necessary when performing dynamic power management in order to calculate the remaining power and make power-up decisions. Current PSE software uses dynamic control to obtain the total ports' consumed power, reading each port's current and voltage and summing the ports' power together.

Figure 2. Mutual identification between PSE and PD

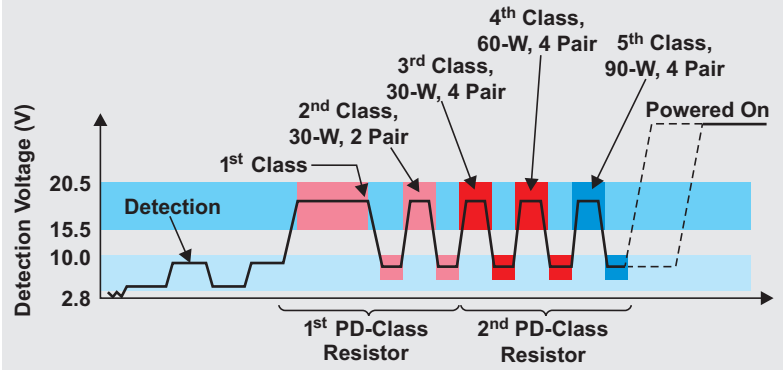

Figure 3. PoE system block diagram

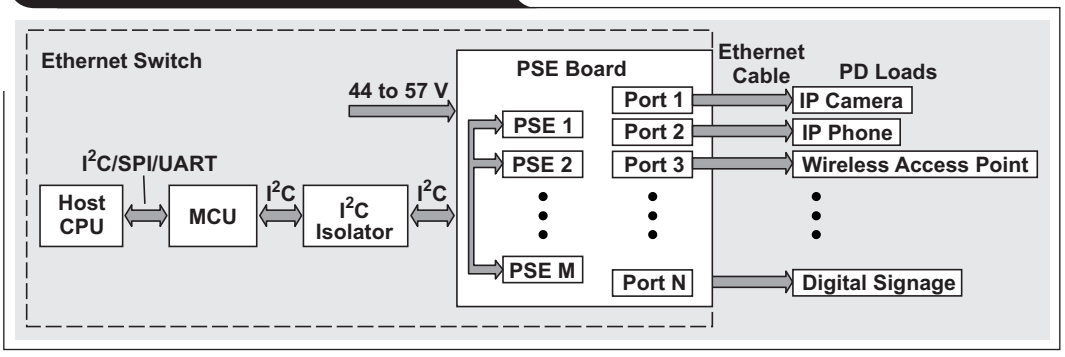

Figure 4. Timing waveform for getting current for one port

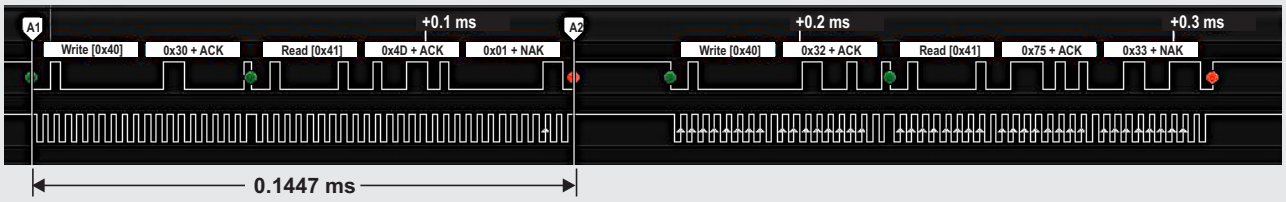

In a system with a high port count, calculating the total consumed power using dynamic control is time consuming. For example, assume that I^2C (400 kHz) is the communications protocol between the PSE and the MCU. Calculating one port's power includes reading two registers for current and two registers for voltage. Reading one port's voltage information requires five I^2C transactions, which takes 1/400 kHz × 9 bits × 5 = 0.11 ms. In total, it takes 0.22 ms to get one port's voltage and current to perform power calculations. Theoretically, it takes 0.22 ms × 96 channels = 21.12 ms to finish one cycle of power calculations.

Figure 4 shows some experimental data captured on an I^2C bus. It took 0.1447 ms to finish one voltage read and another 0.1447 ms to finish one current read. So except for the time that it took to store the data and set up the registers, reading current and voltage from one port took 0.2894 ms. For a 96-channel system, it will take at least 27.7824 ms to finish the reading.

- **Multiple power supplies:** Some systems include multiple AC power supplies. Sometimes, these supplies work in tandem to deliver the required total power, and at other times, the second supply is a redundant backup supply. Both systems need to somehow account for the possibility that one of the two supplies could fail or go offline temporarily. When this happens, there is a need to turn off low-priority ports quickly to avoid power-supply overload.

PSE design solutions

There are solutions to each of the three technical challenges in PSE design.

- **Port power management:** As stated earlier, the goal of port power management is to power up as many PDs as possible while keeping the total power consumption under the power budget and limiting the power cycling of PDs. As shown in Figure 5, port power can be managed by limiting the PD's class level or predefining power limits. There are two ways to calculate the remaining power: dynamic mode or static mode. Dynamic mode uses the actual power consumption to calculate the remaining power, while static mode uses the port power limit. For example, if the total system power budget is 100 W, the PD's class could be limited. With a Class-4 PD, where the maximum power consumption could be 30 W, turning on one port would consume 20 W. In dynamic mode, the remaining power would be 100 W – 20 W = 80 W. In static mode, the remaining power would be 100 W – 30 W = 70 W. There are two ways to prioritize the ports: user-defined priority and first come, first served.
- **Load-step transitions:** It is possible to measure PSE power by using an external sensing circuit at the system's input instead of getting the individual-port's voltage and current readings from the PSE register. The MCU can be used to quickly shut down low-priority ports to prevent power-supply overload, which would improve system response latency. The MCU is continuously reading the analog-to-digital converter (ADC) and monitoring the total consumption of input power. Although the ADC sampling rate determines the system response time, it can be adjusted to meet different requirements.

Figure 5. Power-management methodologies

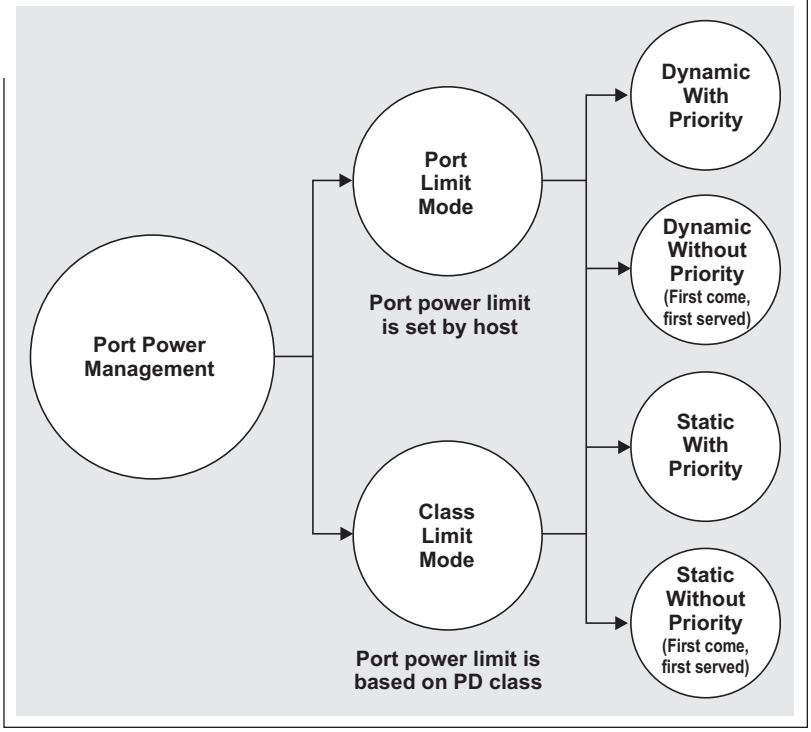

When an overload condition occurs, the MCU will perform the calculation and turn off low-priority ports to maintain input power consumption within the power budget, thus avoiding a system reboot. When the overload condition is removed, the MCU will incrementally power up the ports it shut down if the power budget allows.

Another method of power management uses a comparator to generate a voltage-signal step change when the overload current meets a predefined current threshold. The signal is fed into the general-purpose input/output of the MCU that generates a hardware interrupt and performs a fast port shutdown. Again, when the overload condition is removed, the MCU will incrementally power up the ports it shut down if the power budget allows. This process can reduce system latency from 127 ms to under 1 ms.

- **Multiple power supplies:** One solution is to connect the power-good (PG) signals for the power supplies to the MCU's hardware-interrupt pins. The MCU assigns different fast-shutdown priorities to different ports based on the power budget and that port's priority level. If one power supply malfunctions, the MCU sends predefined fast-shutdown (OSS) signals to the PSE and shuts down low-priority ports to avoid overloading the power supply. Test data showed that the MCU can shut down low-priority ports in less than 200 µs.

  Figure 6 shows a block diagram of a PSE system that requires two power supplies. Each power supply has a power-good indication signal connected to the MCU as an interrupt.

Figure 6. PSE system with multiple power supplies

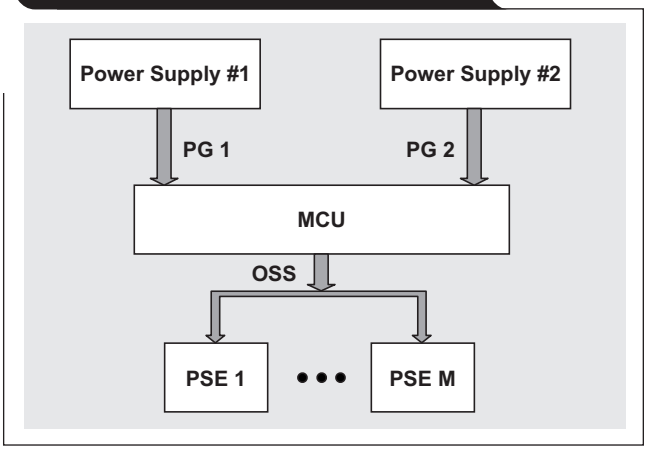

Conclusion

The biggest challenges in PSE design are always related to power management because the power supplies are usually not rated at full power due to size and cost. It can be expensive to solve system-level problems at the integrated-circuit level, but there are more cost-effective solutions through system software.

TI provides PSE software solutions that address the system challenges described in this article. Configuring systems from the host CPU through I^2C or UART dramatically reduces the design cycle and saves resources.

Related Web sites

Product information:
TPS23881
TPS23882

TI PoE solutions

TI Worldwide Technical Support

TI Support

Thank you for your business. Find the answer to your support need or get in touch with our support center at

www.ti.com/support

China: http://www.ti.com.cn/guidedsupport/cn/docs/supporthome.tsp

Japan: http://www.tij.co.jp/guidedsupport/jp/docs/supporthome.tsp

Technical support forums

Search through millions of technical questions and answers at TI's E2E™ Community (engineer-to-engineer) at

e2e.ti.com

China: http://www.deyisupport.com/

Japan: http://e2e.ti.com/group/jp/

TI Training

From technology fundamentals to advanced implementation, we offer on-demand and live training to help bring your next-generation designs to life. Get started now at

training.ti.com

China: http://www.ti.com.cn/general/cn/docs/gencontent.tsp?contentId=71968

Japan: https://training.ti.com/jp

A011617

SLYT781

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI’s products are provided subject to TI’s Terms of Sale (www.ti.com/legal/termsofsale.html) or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI’s provision of these resources does not expand or otherwise alter TI’s applicable warranties or warranty disclaimers for TI products.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265