MARVELL

WHITE PAPER

Power over Ethernet (PoE): An Energy-Efficient Alternative

Roman Kleinerman, Marvell
and Daniel Feldman, Microsemi Corporation

May 2011

Introduction: How Ethernet Switches Consume Power

In modern networking infrastructure, information technology (IT) and data center managers are seeking "green" alternatives for reducing the power consumed by Ethernet switch/router and server equipment. Such reduction is required to create environmentally friendly products and significantly reduce the operational cost. For example, in 1993 internet traffic for the entire year totaled a few hundred terabytes. In 2010, 17 years later, internet traffic accounts for this much data *per second*! In fact, today more than 50 percent of a data center's operating expense is spent on cooling, for enabling on-board fans and for the air conditioning.

Traditionally, network equipment has been designed and for highest performance without clear metrics for power consumption and energy efficiency. Specifically, power efficiency has been decoupled from networking devices supporting the Power over Ethernet (PoE) protocol. Consequently, the networking market has experienced a rapid increase in power consumption, particularly with using high-GHz processors.

Power losses due to idle circuitry present a big and prevalent concern, taking into consideration the more than 300 million Ethernet switch ports shipped every year. The Institute of Electrical and Electronics Engineers (IEEE) Energy Efficient Ethernet specification was defined to significantly reduce the power consumption of the 600 million-plus Ethernet ports shipped per year. However, this specification does not address cases when PoE is deployed, where the majority of the power losses occur in the PoE delivery subsystem—not in the data portion.

With nearly 70 million PoE switch ports expected to ship in 2010, this is a major concern for companies deploying IP telephony, WLAN networks, IP security and other applications powered over Ethernet infrastructure. For example, a typical 48-port Ethernet switch has a 50W to 80W power supply allocated for the traditional Ethernet switch and transceiver IC, with a 370W to 740W power supply allocated solely for PoE. This is a factor of approximately 8-to-1, which means that even minor gains in PoE efficiency may greatly improve the overall efficiency of an Ethernet switch.

Traditional Energy Efficient Ethernet

To address the increasing power consumption of Ethernet switches, the IEEE developed and ratified the 802.3az standard—called the Energy-Efficient Ethernet (EEE) standard—that implements low-power idle (LPI) modes for Ethernet Base-T transceivers (100Mb, 1GbE and 10GbE), as well as the backplane physical layer.

The primary concept behind the EEE standard is to power links down in periods of low utilization, or when completely idle, and then power links back up when they are transmitting data again. This concept is based on the well understood fact that typical client/server Ethernet links in a typical network environment are idle most of the time, while traffic bursts occur only occasionally.

The EEE standard, which specifies the protocol of LPI signaling for both sides of the physical link, enables rapid adjustments of power-saving modes for connected devices—including powering down transmitting and receiving functions while no data to sent. In addition, EEE defines a protocol that enables Ethernet PHYs in LPI mode to keep operational parameters updated, thereby preserving link stability and avoiding disconnections. Furthermore, EEE specifies a definition of signaling by one of the sides that indicates when the physical link is needed and enables fast link "wake up." As a result, in state-of-art Ethernet PHY technology, LPI could save up to 1W per Ethernet link. Still, EEE does not address PoE power consumption and how that could be reduced.

PoE as a Power-Saving Mechanism

One of the compelling reasons for utilizing PoE to power devices, rather than traditional alternating current (AC) supply bricks, is the ability to remotely shut down devices—not to mention reducing cable installment. By controlling when devices are on and off, a significant amount of power can be saved. For example: nighttime cameras can be shutdown during the day (and vice versa) from a centralized point; IEEE 802.11 WLAN access points can be turned on to increase coverage/bandwidth, or turned off at times of low utilization; and IP phones can be turned off during evenings, weekends or periods of non-use.

In multi-port installations, statistics also demonstrate PoE's advantage. While a single AC power brick must support a device in all of its operating modes, a shared power supply for multiple PoE devices can be sized according to the average power utilization—in the same manner as POTS telephony has been employed for years. This represents a major reduction of idle switching power supply losses, which typically account for 10- 20 percent of the maximum power supply load. And when more power is necessary, it is possible to add an additional power supply to a PoE switch or midspan, ensuring that the power is sized according to the growing needs of a business.

Energy Efficient PoE: Advanced PoE Power-Saving Techniques

With the evolution of PoE from a fairly low power source (up to 12.95W per port) to one with devices of up to 25.5W, the direct current (DC) power losses over Ethernet cables increased exponentially. Approximately 4.5W/port of power is wasted on a CAT5, CAT5e, CAT6 or CAT6A cable, all with 25ohm worst-case loop resistance after 100m (i.e., cable transmission efficiency is only 25.5/30, equaling 85% worst case). Even when using a voltage of 54V instead of the 50V worst case, efficiency does not exceed 87%. Remember, EEE typically saves no more than 1W per link, so addressing the 4.5W per link loss from PoE transmission inefficiency would provide much more incremental savings.

New energy-efficient PoE (EEPoE) technology can change increase efficiency to 94% while transmitting over the same 25ohm cable, powering IEEE 802.3at-compliant devices in synchronous 4-pairs. When utilizing synchronous 4-pairs, powered devices are fed using all the available wires. For example, on a 24-port IEEE 802.3at-2009 Type 2 system (delivering 25.5W per port), more than 50W are saved.

Specifically, reference switch systems with both EEE and EE PoE built into products demonstrate this significant savings for the industry, as demonstrated in the table below:

Technology	How much it could save	When will the saving kick-in	Number of Ethernet Links shipped per year
EEE over GbE	~1W per link	Power saving is instantaneously when link is idle	200M and growing
EEPoE with .3af	~0.6W per link	Power saving will occur as long as power is on	40M and growing
EEPoE with .3at	~2.1W per link	Power saving will occur as long as power is on	30M and growing
Total Savings for EEE + EEPoE	***3.13W per link***		

Note: EEPoE technology does not require changes on the powered device side, thereby providing immediate saving just by upgrading the switch or midspan.

Applying Innovative Technology to PoE:

To help advance EEPoE technology, Microsemi® and Marvell® leveraged the unique microprocessor technology embedded in Marvell® Prestera® DX4100 and DX2100 product families. This innovative approach equips customers with a platform for development of cost- efficient yet powerful network devices with the built-in option of main CPU offloading. Along with an integrated main CPU, the Marvell Prestera DX4100 and DX2100 microprocessors are equipped with an additional embedded Marvell® Dragonite microcontroller with the ability to run light-weight, stand-alone applications. Today, Microsemi EEPoE management software can run on the Dragonite microcontroller, enabling customers to increase PoE performance while reducing overall PoE system cost by 15 percent.

Summary

With the ratification of the IEEE 802.3az standard in September 2009, Power over Ethernet (PoE) is ready to deliver further power savings. PoE technology, which enables devices to be powered over the same Ethernet cabling infrastructure as data, provides important benefits. It eliminates costly AC outlet installation and allows devices to be deployed without cable installation, while automatically preserving power based on whether a device is on, off or transmitting data.

Combined with the EEE standard, PoE is poised to deliver even greater benefits. The two should not be decoupled, but rather be leveraged in complementary fashion. Technology providers such as Microsemi and Marvell are working together to develop innovative energy-efficient PoE solutions (e.g., switches, midspans) that enable OEMs to offer more cost-effective and environment-friendly products—something the entire industry will benefit from in both the near- and long-term.

About the Authors:

Roman Kleinerman
Marketing and Product Line Manager, Marvell

Roman Kleinerman is marketing manager and product line manager at Marvell, responsible for the Marvell Prestera DX family of packet processors and Avanta™ family of Marvell UPON (GPON/EPON and Active Ethernet) products. Previously, Mr. Kleinerman worked as software team leader and project manager for Radlan Computer Communications, specializing in real-time embedded programming. He holds B.S. degrees in mathematics and computer science from Negev University in Israel.

Daniel Feldman
Director of Marketing, Microsemi's AMS Group

Daniel Feldman is director of marketing for Microsemi's AMS Group, responsible for the profit and loss (P&L) of all products in the PoE, Ring Generation and Solar markets, including ICs, modules and midspans. He was a member of the IEEE 802.3at Task Force and is chairman of the Ethernet Alliance PoE/PoEPlus Technical Committee. Previously, Mr. Feldman worked as senior product manager for PowerDsine, as System Architecture Group manager for IC4IC, as VHDL engineer for NICE Systems and as VLSI engineer for RAFAEL. Mr. Feldman holds a B.S. degree (Cum Laude) in computer engineering from the Technion in Israel.

Marvell Semiconductor, Inc.
5488 Marvell Lane
Santa Clara, CA 95054, USA

Tel: 1.408.222.2500
www.marvell.com

Power_over_Ethernet-001_white_paper 5/2011