Energy Efficient (Power over) Ethernet

April 2011

Authors:
Daniel Feldman, Microsemi
Roman Kleinerman, Marvell

Introduction

With over 300M Ethernet ports shipped worldwide, the networking industry is effortlessly working on standards and technologies to reduce energy consumption. This paper describes the benefits of employing Energy Efficient Ethernet (EEE) as specified in IEEE802.3az-2010 and Energy Efficient Power over Ethernet (EEPoE) as allowed by the IEEE802.3at-2009 standard to reduce the energy consumption of Power over Ethernet Switches.

How Ethernet Switches Consumer Power

In modern Networking Infrastructure the IT and data center managers are looking for green alternatives targeted to reduce the power consumed by Ethernet Switch/Router and Server equipment. Such a reduction is a key for creation of environmentally friendly products as well as significantly reducing the operational cost. For an instance, in 1993 total Internet traffic for the entire year was few hundred terabytes; in 2010, 17 years later - this is the amount of traffic per second only! Today more than 50% data center's OPEX is actually spent on cooling, both for enabling on-board fans as well as for the air conditioning.

Traditionally, network equipment has been built with design considerations for highest performance without clear metrics for power consumption and energy efficiency. Particularly power efficiency decoupled from networking niche of devices supporting Power over Ethernet (PoE) protocol. As a consequence, the networking market witnessed a rapid increase in power consumption, particularly with using high-GHz processors.

Power losses due to idle circuitry are a big and prevalent concern in view of the over 300M Ethernet switch ports shipped every year. IEEE Energy Efficient Ethernet specification was defined to significantly reduce the power consumption of the 600M+ Ethernet ports shipped per year. However, this does not address the cases when Power over Ethernet is deployed where the majority of the power losses occur in the Power over Ethernet delivery subsystem and not in the data portion. With almost 70M PoE switch ports expected to ship in 2010, this is now a major concern for companies deploying IP telephony, WLAN Networks, IP security, and other applications powered over an 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 around 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, the IEEE came out with 802.3az standard, a.k.a. 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 main EEE idea is to power links down in periods of low utilization (or completely idle) and to power links back up when they need to transmit data again. This idea based on a well known fact that typical server/client Ethernet link in typical network environment is idle most of the time while traffic bursts occur only occasionally.

The IEEE 802.3az specifies the protocol of LPI signaling for both sides of the physical link, enables rapid adjustments of power saving modes of the connected devices, and allows powering down transmitting and receiving functions while no data is sent. In addition, EEE defines a protocol that enables Ethernet PHYs in LPI mode to keep operational parameters up to date preserving link stability and avoiding disconnections. EEE defines a way of signaling by one of the sides, that the physical link is to be used and allows fast link “wake up”. In the state-of-art Ethernet PHY technology, LPI could save up to 1W per Ethernet link yet EEE does not address the PoE power efficiency and how that could be reduced.

PoE as a Power Savings Mechanism

A major reason for utilizing PoE to power devices instead of using AC supply bricks is the capability to remotely shutdown devices and to reduce cable installment. By controlling when devices are on and off, a dramatic amount of power can be saved. For example, nighttime cameras can be shutdown during the day (and vice-versa) from a centralized point, IEEE802.11 WLAN access points can be turned on to increase coverage/bandwidth or off at times of low utilization, and IP phones can be turned off during evenings or weekends. In multi-port installations, statistics also play in PoE’s favor. While a single AC power brick needs to support a device in all of its operating modes with a shared power supply by multiple POE devices, the power supply can be sized according to the average power utilization in the same manner employed for years in POTS (Plain Old Telephone System) telephony. This represents a great reduction of idle switching power supply losses which typically amount to 10% to 20% of the maximum power supply load. When more power is necessary, it is in many cases possible to add an additional power supply to a PoE switch or Midspan while making sure that the power is sized according to the growing necessities of a business.

EEPoE: Advanced PoE Power Saving Techniques

With the evolution of PoE from a fairly low power source (up to 12.95W per port) to one where devices of up to 25.5W are deployed, the DC power losses over the Ethernet cables increased dramatically. When the worst case power supply is used in a PoE switch, 4.5W/port of power is wasted on a CAT5, CAT5e, CAT6, or CAT6A cable. While this can be improved by increasing the PoE power supply voltage to 54V, cable losses remain very high at 3.62W per port. EEE typically saves no more than 1W per link so addressing the 3.62W to 4.5W losses per link become crucial to create PoE switches that are energy efficient.

Energy Efficient PoE (EEPoE) can reduce cable losses to only 1.55W per link while powering a 25.50W IEEE802.3at-compliant device.. EEPoE is allowed by IEEE802.3at-2009 and is based on delivering power over all 4-pairs in a CAT5 or better cable synchronously. On a 24-port IEEE802.3at-2009 Type 2 (delivering 25.5W per port) system, over 50W are saved.

Relying on the fact that every IEEE802.3af-2003 and IEEE802.3at-2009 has to be able to receive power over all the 4-pairs to be compliant with the standard, EEPoE allows PoE switch manufacturers to reduce the power dissipation to deliver Power over Ethernet to not only new Powered Devices but to the huge installed base of over 100 million powered devices already installed worldwide.

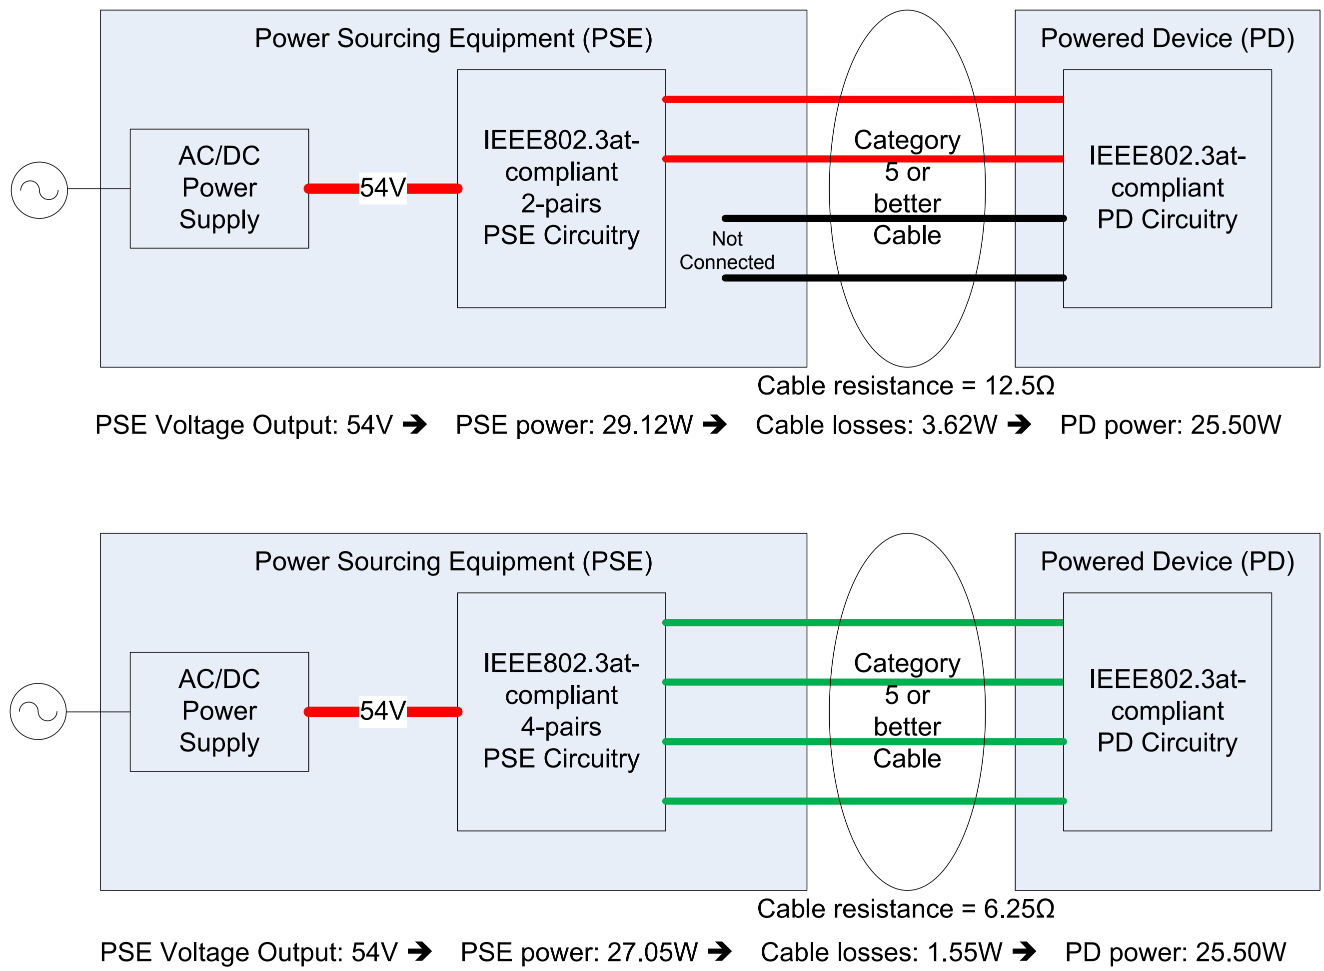

Figure 1: PoE vs. EEPoE Cable losses when powering a 25.50W PD from 54V PSE power supply

Reference Switch Systems with EEE and EEPoE has been built to demonstrate this saving for the industry.

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	Instantaneously, when link is idle	200M and growing
EEPoE with .3af	~0.6W per link	Saving will achieved as long as power is on	40M and growing
EEPoE with .3at	~2.1W per link	Saving will achieved 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, hence providing immediate saving just by upgrading the Switch or mid-span.

Summary

With total savings of up to 3.13W per link and capable of delivering energy savings to all installations employing standard-compliant PoE Powered Devices, EEPoE presents a technology that can dramatically change the energy footprint of networks which employ PoE as a power distribution and management technology.

Biographies

Daniel Feldman serves as a Director of Marketing at Microsemi's AMS Group, in charge of the 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 Chairs the Ethernet Alliance PoE/PoEPlus Technical Committee.

Previously, Mr. Feldman worked for PowerDsine as a Senior Product Manager, at IC4IC as the System Architecture Group Manager, as a VHDL Engineer at NICE Systems and as VLSI Engineer at RAFAEL.

Mr. Feldman holds a B.Sc. (Cum Laude) in Computer Engineering from the Technion in Israel.

Roman Kleinerman is a Marketing Manager and Product Line Manager at Marvell Semiconductor, Inc., in charge of Marvell Prestera DX Family of Packet Processors and Avanta Family of Marvell UPON (GPON/EPON and Active Ethernet) Products.
Previously, Roman Kleinerman worked in Radlan Computer Communications as a SW Team Leader and Project Manager - specializing in RT embedded programming. He holds B.S. degrees in mathematics and computer science from Negev University in Israel.

Mr. Kleinerman is a chess and wine enthusiast, with a strong belief in a winning combination of both.

About Ethernet Alliance

The Ethernet Alliance is a community of Ethernet end users, system and component vendors, industry experts and university and government professionals who are committed to the continued success and expansion of Ethernet. The Ethernet Alliance brings Ethernet standards to life by supporting activities that span from incubation of new Ethernet technologies to interoperability demonstrations, certification and education.