White Paper

Broadcom® Energy Efficient Networking

Broadcom continues to drive innovation in communications technologies through leadership in and contribution to the development of standards. Broadcom has embraced the IEEE 802.3az Energy Efficient Ethernet™ (EEE) standard within its broad framework for implementing high-performance and low-power technologies at an optimal cost point across its wired Ethernet portfolio.

Broadcom's Energy Efficient Networking builds upon the requirements of the standard by adding control policy as well as hardware and software subsystems in a standards-friendly way. Energy Efficient Networking enables Broadcom customers to build complete energy-efficient systems with enhanced energy savings at a faster time-to-market.

January 2012

Traditional Versus Next-Generation Networks

Traditional networks have focused on optimizing performance and reducing operating costs. Broadcom's next-generation networking solutions add optimized power consumption to make the network energy-efficient.

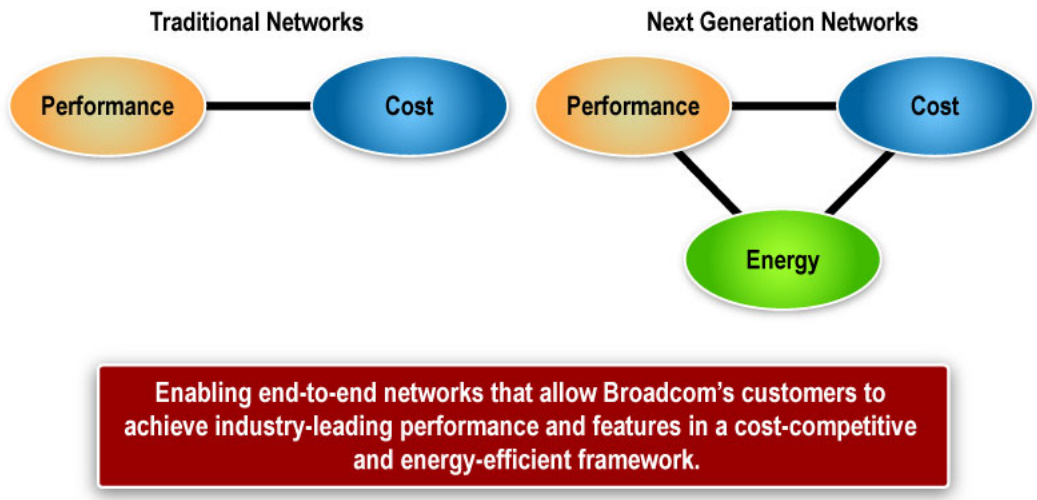

Figure 1: Cost-Performance-Power Model for Traditional Networks vs. Next-Generation Networks

What Is IEEE 802.3az Energy Efficient Ethernet?

IEEE 802.3az, which was ratified on September 30, 2010, is also known as Energy Efficient Ethernet. EEE is targeted at saving energy in Ethernet networks for a select group of PHYs. The PHYs selected in this project include the popular 100BASE-TX and 1000BASE-T PHYs, as well as emerging 10GBASE-T technology and backplane interfaces, such as 10GBASE-KR. The method of power savings currently planned for these PHYs is a technique known as Low Power Idle (LPI).

The legacy Ethernet standards for interfaces of 100 Mbps and higher have an idle state. The idle state requires the bulk of the circuitry to remain powered up, irrespective of data transmission. This condition results in comparable power consumption, regardless of whether there is data on the link. LPI provides for a lower-consumption energy state that can be employed during periods of low link utilization (high idle time), which is common in many Ethernet networks. LPI allows for rapid transitions back to the active state for high-performance data transmission.

Wide Applicability of EEE

Ethernet has become the ubiquitous technology of choice for wired connectivity. Indeed, Ethernet is being deployed in enterprise, small and medium business (SMB), service provider, and home networks as well as professional AV networks. It is also prevalent in data centers and storage applications. As such, all areas of the network stand to benefit from EEE energy savings.

High Potential for Energy Savings

Typical Ethernet traffic profiles have characteristically low average link utilization over time, with occasional bursts that are associated with network activity. This lends itself very well to EEE, which takes advantage of the high percentage of idle time on the link (see Figure 2).

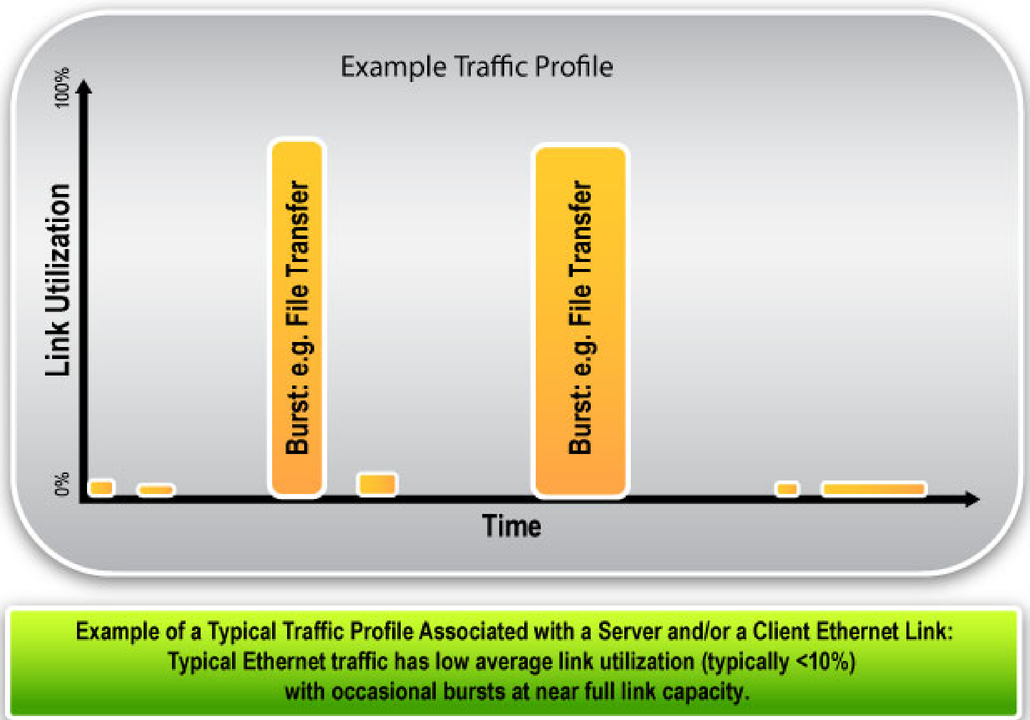

Figure 2: Typical Traffic Profile Example

Ethernet traffic profiles vary based on applications and the market segment. By using EEE in conjunction with the Broadcom control policy technology described later, end customers can take advantage of the network idle times to realize high levels of power savings.

Broadcom Energy Efficient Networking — End-to-End Energy-Efficient Framework and Savings

Broadcom's emerging physical layer products reduce energy consumption. In addition, the controller and switch products allow for additional savings, which can be accomplished only by using EEE LPI beyond the physical layer.

A Framework for "Green" Networking

The following equation illustrates the factors that affect the cost of energy and where the potential for savings can occur.

$$\mathbf{E_T = (P_{active} \times T_{active}) + (P_{idle} \times T_{idle})}$$

$\mathbf{E_T}$ represents the total amount of energy consumed over a period of time, **T**. If you think of your monthly energy bill, $\mathbf{E_T}$ is the amount of energy use recorded by the meter during that month, **T** is one month, and **P** is the unit rate of energy consumption. Because the cost of energy is directly proportional to the rate of energy consumption over time, the less energy used, the lower your monthly energy bill.

The equation is based on the networking device having two states — an active state and an idle state. The total energy consumed is the sum of the energy consumed while the device is in each state. The amount of energy consumed in any state is the product of the average power consumption while the device is in that state and the length of time that the device is in that state.

To minimize the power consumption[1] or, to be more precise, to maximize the power savings, the following three options are available:

- Minimize $\mathbf{P}_{active}$
- Minimize $\mathbf{P}_{idle}$
- Maximize $\mathbf{T}_{idle}$[2]

To relate the equation to EEE, consider what EEE provides to networking systems. Unlike a stand-alone system, networking systems communicate with their link partners. EEE fundamentally provides a method to move a networked system from an active state to an idle state without compromising network connectivity or remarkably impacting upper layer protocols. In essence, it provides the second set of terms ($\mathbf{P}_{idle} \times \mathbf{T}_{idle}$) in the equation.

1. Of course, this equation can be expanded to take into account a more complex system with multiple energy states for the device. In that case, the power consumption would be the weighted average of the power states multiplied by the time the device spends in that state. Nonetheless, the principles of conservation remain the same.
2. Because the total time is the sum of the idle time and active time, maximizing idle time is the equivalent of minimizing the active time.

Broadcom's Innovations Described Within the Framework

There are two additional components that are essential to building EEE systems and networks that are beyond the scope of the standard but are critical to EEE-enabled devices:

- EEE Control Policy: Provides controls when the physical layer enters and exits the low-power state and is outside the scope of standard. The level of integration of the control policy decision engine with the controls to the physical layer will affect the overall efficiency attained. Moreover, the control policy maximizes savings by maximizing the time the system spends in $\mathbf{T}_{idle}$, while minimizing any performance impact on the network.
- Enhanced Savings: Provides enhanced savings in the device that extend beyond the physical layer at either the transmitting link partner or the receiving link partner (see Figure 3). This savings effect involves reducing both $\mathbf{P}_{active}$ and $\mathbf{P}_{idle}$.

Figure 3 represents an edge device (e.g., server or client) connected to a node (such as a switch) within a network. The devices are connected via an EEE-enabled Ethernet link. The figure breaks each link partner into its significant subsystems in an OSI-like fashion, starting with the physical layer at the bottom, beginning with PHY, and moving up through the stack. Figure 3 also illustrates the additional power savings enabled through Broadcom's standards-friendly Energy Efficient Networking enhancements. The principles also apply in a switch-to-switch connection.

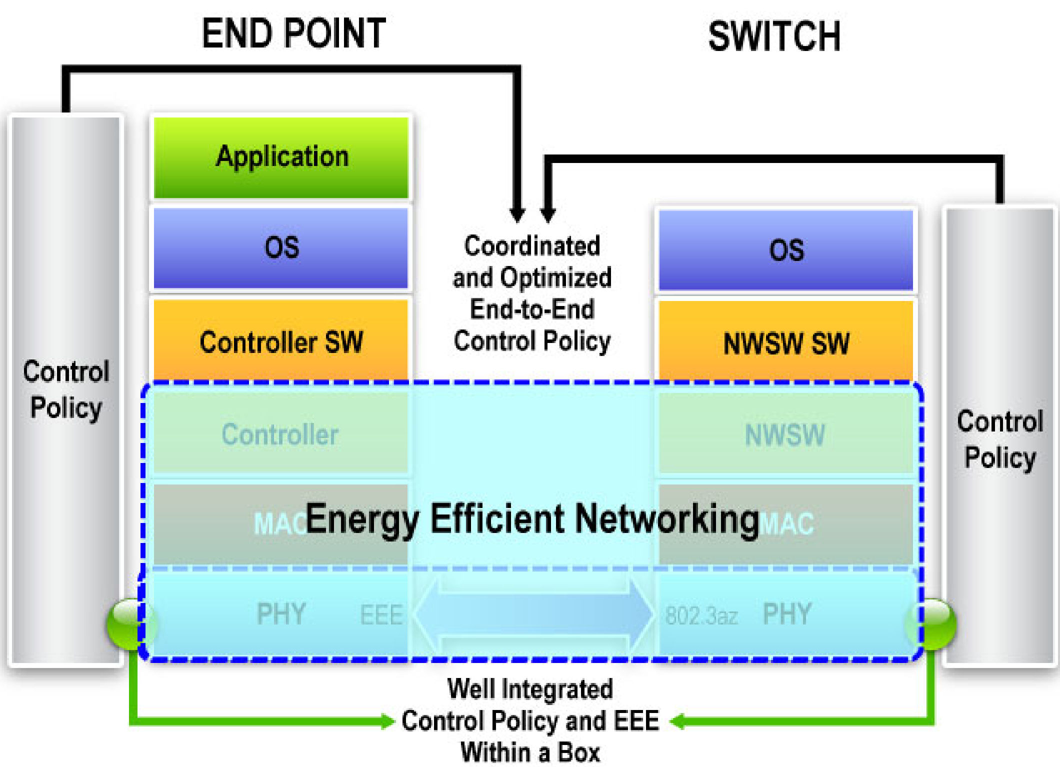

Figure 3: Additional Power Savings Enabled by Broadcom Energy Efficient Networking

Broadcom Energy Efficient Networking — Control Policy Approach

The control policy spans more than just one layer because the decision to go into and out of the energy-saving state will require the management system of the device to act based upon a number of inputs.

The control policy must be customized for the particular application to yield optimal savings, and this can be done with Broadcom software. In addition, the Broadcom software stack allows for programmability for different levels of performance and energy savings options.

Introducing AutogrEEEn™ Technology

To allow for rapid market adoption and immediate customer transition to EEE-enabled ports, Broadcom has introduced AutogrEEEn™ technology as a part of its EEE PHY lineup. AutogrEEEn technology enables a device with a non-EEE MAC to seamlessly transition to EEE capabilities by implementing control policy assist engines and circuitry inside the PHY device.

EEE requires control for the PHY completion via in-band signaling over the MAC/PHY interface. This requires a change to both the PHY and MAC silicon. A number of systems have the MAC and PHY as two different pieces of silicon with the MAC often embedded in a switching or controller-type device. These MAC-containing devices have associated drivers and software and are often multiport devices. Therefore, a transition to EEE may be hampered by additional development that involves replacement of the MAC-containing devices.

Broadcom's AutogrEEEn technology eliminates the need to change the MAC/PHY interface on the MAC silicon and allows for rapid transition today with legacy non-EEE MAC silicon attached to Broadcom AutogrEEEn-enabled PHYs.

Conclusion

This paper introduced Broadcom Energy Efficient Networking, which builds upon the IEEE 802.3az EEE standard in the following ways:

- Using a comprehensive control policy.
- Realizing additional energy savings — putting additional resources "to sleep" beyond the PHY by taking advantage of the EEE low-power state.
- Using software to maximize energy savings in various applications and spaces, allowing for optimized energy savings and a control policy that can be customized.

About the Author

Wael William Diab serves as Senior Technical Director in the Office of the CTO at Broadcom Corporation. In this role, Diab is responsible for defining the technical strategy for the Infrastructure & Networking Group (ING). Prior to Broadcom, Diab served at Cisco in various technical, architectural, and business leadership roles, focused on next-generation networking products and technologies. Diab holds BS and MS degrees in Electrical Engineering from Stanford University, a BA degree in Economics from Stanford, and an MBA with honors from the Wharton School of Business. He has developed over 300 patents and patents-pending in the networking space, with over 70 issued patents in the United States.

Diab also serves as Senior Member of the IEEE and was unanimously elected and reelected as the Vice-Chair of the IEEE 802.3 Ethernet Working Group, serving in that position since 2008. Diab is a member of the IEEE-SA Standards Board, which oversees all IEEE standards activities. He is also a member of the IEEE Standards Education Committee (SEC), was elected to the IEEE-SA Corporate Advisory Group (CAG) and serves as the IEEE-MGA liaison to the IEEE-SA.

Named winner of the 2011 TechAmerica Innovator Awards for his leadership in Green Technology, Diab was recognized by the David Packard Medal of Achievement and Innovator Awards for his leadership in the development of Broadcom's Energy Efficient Networking solutions.

BROADCOM CORPORATION
5300 California Avenue
Irvine, CA 92617

Phone: 949-926-5000
Fax: 949-926-5203
E-mail: info@broadcom.com
Web: www.broadcom.com

EEE-WP102-R January 2012