



## Power Management for Networking Devices

---

---

## Table of contents

Introduction and Background .....	3
Monitoring, Reporting, and Responding to Energy Usage .....	5
Key Power Management Features for Networking Devices .....	8
Key Power Management features for Attached Devices .....	9
Optimizing Network Energy with Application Requirements .....	11
Network Load .....	11
Examples .....	11
Application Requirements and Power States of the Device .....	12
Example .....	12
Other Optimizations .....	13
Detailed Management Requirements .....	13
Conclusion .....	15
References .....	16

## Introduction and Background

As business sectors and consumers shift toward smarter, more connected economies and lifestyles respectively, the demand for computing and networking - in the data center and across enterprise networks to all connected devices - will continue to grow. With this growth in capacity comes a potential for unsustainable energy use growth, as well; the information and communications technology (ICT) industry must therefore address energy consumption through increased energy efficiency of networking equipment and connected devices. Already, one example of positive technology impact can be found in computing energy use in data centers, where newer components and systems plus consolidation of workloads through server virtualization has driven projected energy growth from 100% (EPA, 2007) down to 1-2% over the last five years (Kooimey, 2011). Today, energy consumption in telecommunications and networking faces a similar challenge. Without explicit focus on energy efficiency, the increasingly connected world may face greenhouse gas (GHG) emissions growth by as much as 6% per year (LBNL, 2010). On the operational side of business, unless today's CFOs and CIOs move to adopt best practices for energy efficiency in networks and across connected devices, they will continue to see the negative effect of unneeded energy use on operating costs and productivity.

Launched in 2007, Climate Savers Computing Initiative's (CSCI) original mission was to reduce greenhouse gas emissions from PCs and servers by promoting aggressive improvements in their energy efficiency. In 2010, several leading companies in the networking industry (Broadcom, Cisco, Emerson Network Power, F5, and Juniper Networks) joined CSCI to address network equipment energy consumption and associated greenhouse gas emissions and to establish best practices for the enterprise and telecommunication networking equipment segments. Our approach to reducing energy consumption for network devices is two-fold -- first, accelerate industry's design and product offerings of energy efficient network equipment. Second, educate IT purchasers and network operators regarding best practices and energy efficient considerations when developing and operating network systems for their IT infrastructures. We make a strong case for establishing an open, energy-proportional network device ecosystem that will improve the sustainability of the world telecommunication infrastructure.

To address different facets of the energy consumption opportunities within network equipment, CSCI is launching a series of white papers. These white papers are divided into three distinct but interconnected topics:

1. Considerations for Selecting Power Supplies for Networking Equipment and Evaluating Power Conversion Efficiency
2. Power Management for Networking Devices
3. Energy Efficiency Guide for Networking Devices

Together these three white papers will direct system designers, IT managers, and IT procurement professionals to design, purchase, install, and manage energy efficient networks.

This paper addresses the second topic, device power management and control. Specifically, it describes the importance of monitoring the power and performance characteristics of networking devices and communicating device-level information to network management systems. It then examines system and device design opportunities for improved power savings, laying out key power management (PM) features for networking devices and devices attached thereto. Next, it considers how application behavior and resource-use interact with energy-use in networking devices. Finally, it presents a summary of device design requirements pertinent to PM needs. System design engineers will be able to use the information in this paper as a concrete guideline for addressing power and performance measurement capabilities in networking systems. IT & Infrastructure Managers will be able to understand energy consumption in networking equipment in order to procure networking equipment that provides better power savings and to design and operate an optimal network, hence reducing GHG emissions. Examples of Power Management (PM) opportunities focus on Energy Efficient Ethernet (EEE)<sup>1</sup> and Power over Ethernet (PoE)<sup>2</sup> as non-proprietary, standards-based supports for PM-based energy use curtailment.

Contributors to this paper include the following members of the Climate Savers Computing Networking Workgroup. Views expressed in this paper represent the opinions of the authors and do not necessarily represent the views of their affiliated organizations or companies.

Hugh Barrass	Cisco
Akella Chakravarthy	Intel
Pierre Delforge	Natural Resources Defense Council
Wael Diab,	Broadcom Corporation
Kathleen Fiehrer	Intel
George O. Goodman	Climate Savers Computing Initiative
Shaun Harris	Microsoft
Steve Haughey	Juniper Networks
Rajiv Limaye	Cisco
Charlie Swiontek	PowerOne
Oliver Tavakoli	Juniper Networks

---

<sup>1</sup> First published as IEEE Std. 802.3az-2010 <http://standards.ieee.org/about/get/>

<sup>2</sup> First published as IEEE Std. 802.3af-2003 and IEEE Std. 802.3at-2009, now incorporated into IEEE Std. 802.3-2008 <http://standards.ieee.org/about/get/>

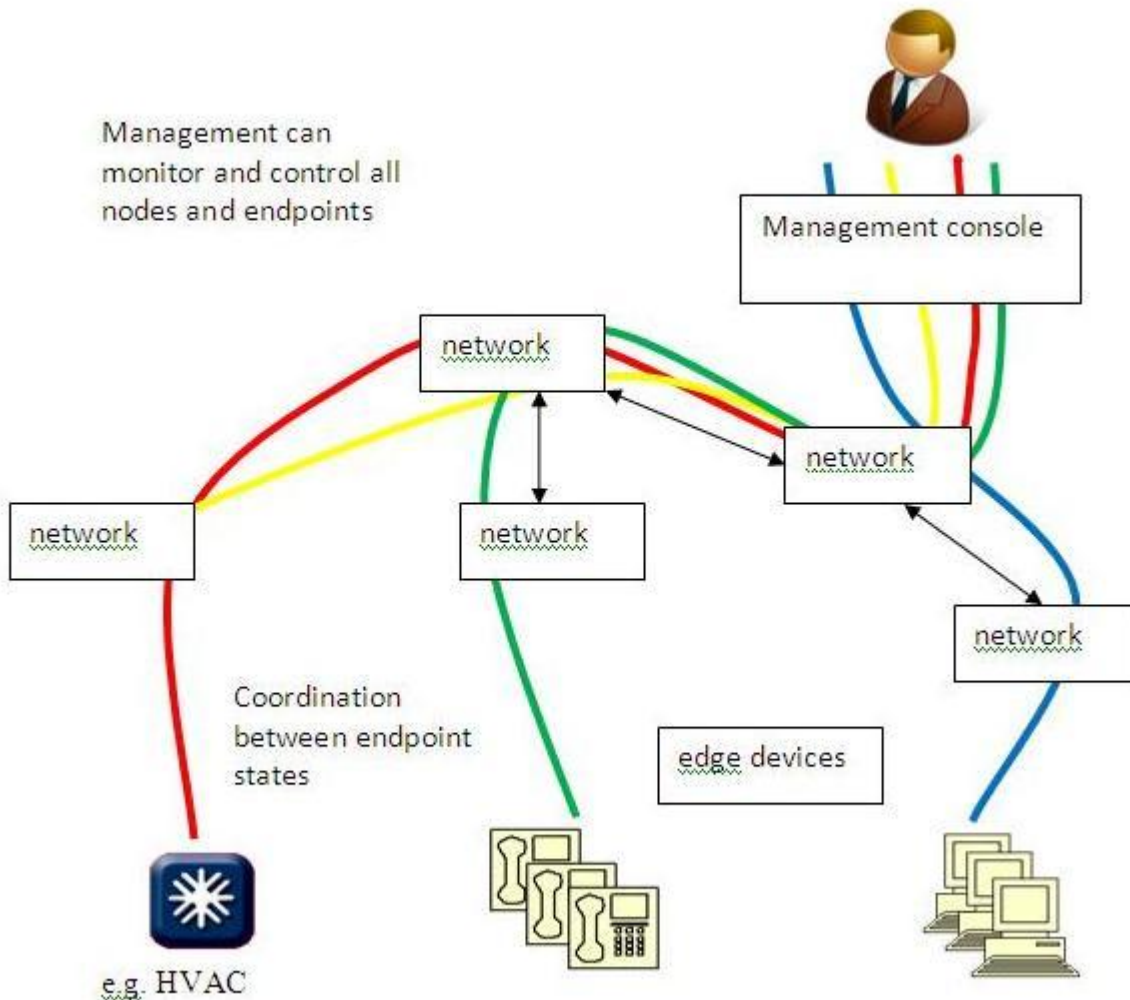
## Monitoring, Reporting, and Managing Energy Usage

Monitoring power usage by individual network devices in real time enables effective management of the network. Accurate power information is an important enabler to forecasting demand, controlling cost, balancing load on infrastructure devices, and planning for changes in the network. Understanding the power consumption of the network enables optimization and can return real savings on energy to meet financial and eco-efficiency objectives without compromising network performance.

Measuring, monitoring and responding to system energy use are key to optimizing the overall energy performance of network systems. Mapping the components of system power consumption, including input power, used power, total power capacity, system throughput, and system utilization measured over time, can identify areas for energy saving opportunities. However system-level decision-making, looking at the complete picture rather than just components, regarding power management allows for further energy efficiency. Such a system-level view requires knowledge of both power usage and workload data in real time for individual devices and, through management systems, for complete networks; these individual and aggregated data can enable PM actions to achieve energy optimization. Consider some examples:

- Identifying usage patterns allows management programs to shut down unused system
- Monitoring power supply utilization allows network systems to manage redundant Power Supply Units (PSUs), eliminating needless energy consumption during low efficiency operations.
- Managing fan speeds in relation to power delivery also saves energy and therefore costs.
- Monitoring system energy efficiency, network owners can determine whether a system is consuming power proportionate to work being done. If a system is not performing proportional to power consumption, network design or system/component replacement is an option to improve energy efficiency.
- Devices which support 'sleep' modes in communications with their peers when there is no data to transmit should be configured to achieve maximum power savings while still meeting the business requirements for the network.

In order to accomplish the monitoring, reporting, and management approaches described above, devices must have some specific PM features. Those features – for networking devices and those devices attached thereto – are the topic of the following sections.

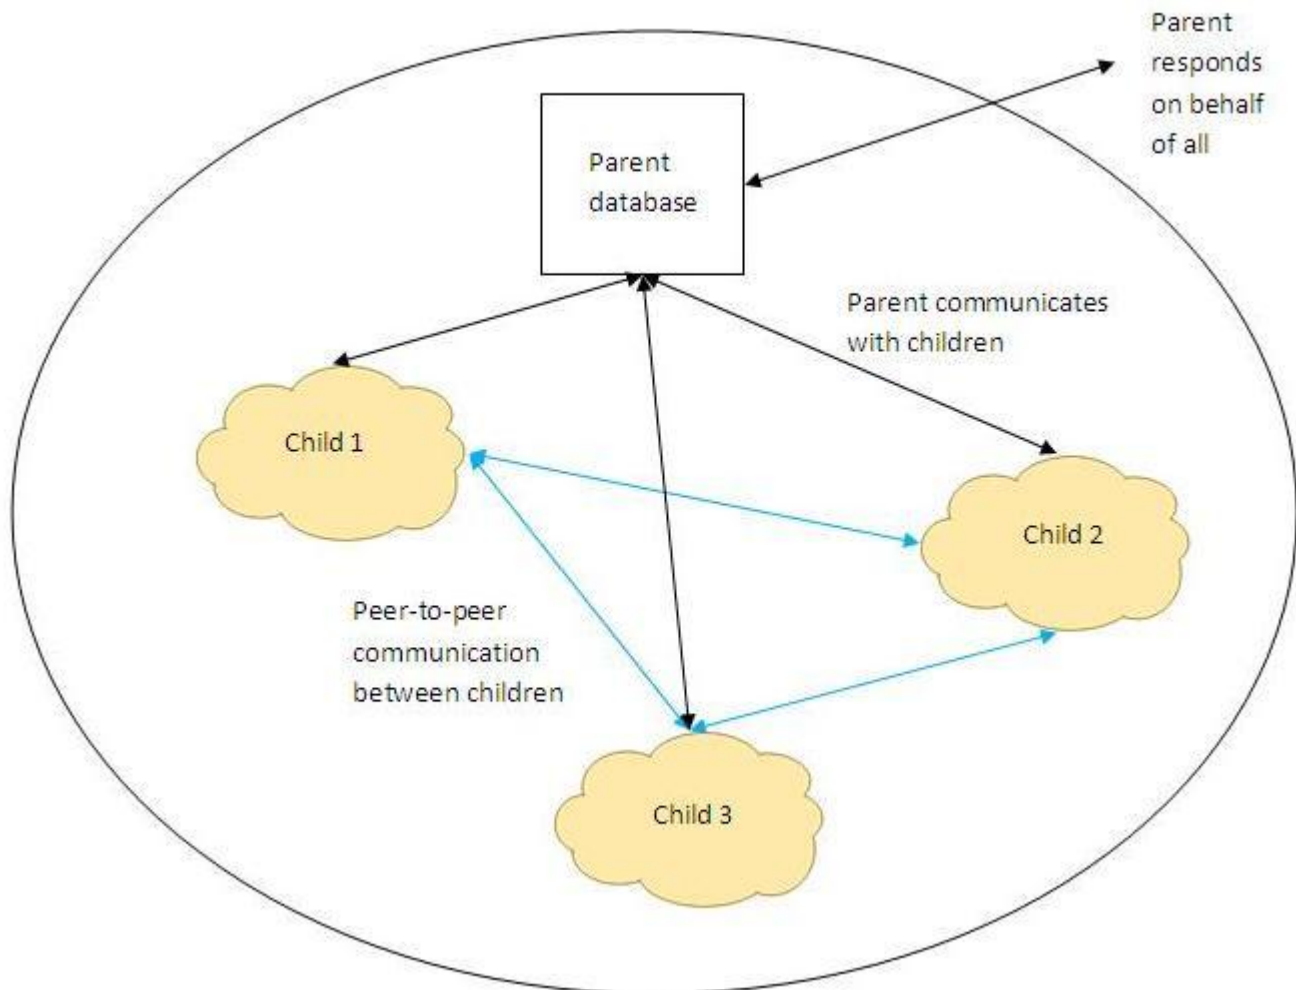


**Figure 1: Measuring, monitoring and responding to system energy use are key to optimizing the overall energy performance of network systems.**

**Coverage and granularity of monitoring and reporting:**

Networks with a high percentage of devices that support power monitoring provide the best opportunity to establish an energy efficient infrastructure. Managed devices in the network should, at a minimum, measure and report the cumulative energy use of the device (e.g. KWhr or Joules). Network devices that have redundant power supply subsystems should include power levels from each line source as well as total power. Network devices that directly provide power for external or attached devices, through PoE for example, should measure and report energy usage of those external devices. Some equipment in the network may not support power management (e.g. low end devices, non-managed elements or older equipment); this equipment can be monitored via external smart Power Distribution Units (PDUs) that may report line power via Simple Network Management Protocol (SNMP) and thereby extend the visibility of those network elements.

In some cases there may be too many manageable instances to deal with efficiently. In this case, the management infrastructure may make use of a “parent and child” architecture in the network as being defined by IETF Energy Management working group<sup>3</sup>. This architecture allows networking nodes to collate data from attached devices and present it as a summary to the management interface. There are also definitions for interrogating and controlling devices in the network according to device type or power priority. For example, the management system may ask for information regarding all IP phones attached to the network, or may turn off all desk lights in a region. The power priority feature may be used to cut power to non-essential devices during a power outage or supply crisis.



**Figure 2: Parent and child management architecture.**

<sup>3</sup> <http://datatracker.ietf.org/wg/eman/>

### **Power Reporting and Management example using Power over Ethernet (PoE):**

The MIB for PoE is defined IEEE Std. 802.3.1<sup>4</sup>. It should be noted that the PoE specific managed objects in the MIB (prefixed with “peth”) offer only a rudimentary control of the function, whereas the detailed control and monitoring is achieved using the LLDP objects (prefixed with “lldp” – for devices that support sophisticated PoE management). Currently there is no standard definition for management objects that reflect the actual power usage, either from the perspective of the Power Sourcing Equipment (PSE) or the Powered Device (PD). It is recommended that PSEs should, at a minimum, report the cumulative energy usage of each attached PD (expressed in KWhr or Joules). Appendix A contains a detailed list of managed objects for power management.

For example a PSE device may be supplying power to a number of IP phones (some of which are required to support emergency access) and some advertising display screens. Using SNMP, the energy management system may interrogate the PSE system to find the power that is being used at the moment, as well as the cumulative energy usage (which will be reflected as part of the energy billing for the facility). Additionally, using the IETF energy management framework, the manager may find the power being used by each class of attached devices. During an energy supply crisis, the bulk energy provider might request that the network operator should reduce its power load for a period. The manager can then instruct the PSE to renegotiate the power states for the attached PDs so that low priority devices might be denied power. A further crisis might be caused by a failure of the mains supply for the PSE and a reliance on battery-backed UPS. During this condition, the PSE would remove power from all non-essential PDs and may also renegotiate the power state with the remaining PDs. For example, the IP phones may be instructed to allow emergency calls only or to deny video service.

## **Key Power Management Features for Networking Devices**

Networking devices are generally manageable via their network interfaces using either a management console or a web browser based interface. Within enterprise networking, it is expected that devices should communicate with a network management system using SNMP. To ensure that network managers can understand both the power and performance impacts of any management actions, the devices should minimally support the traffic management and reporting MIBs defined in IEEE 802.3.1-2011<sup>5</sup>.

Energy measurement and control is currently the subject of study within the IETF Energy Management Working Group. The draft framework indicates that a standard MIB will shortly be defined that reports cumulative energy usage (e.g. a running count of Kilowatt-Hours), a measurement of actual power and a record of the highest peak power since the previous reading. Devices will also be able to report the maximum expected power usage for the current operating mode and, if power modes are supported, may be able to report the expected power usage for adjacent available modes. Equipment with multiple power supplies should report the state of each power supply as a separate subsystem.

---

<sup>4</sup> First published as IEEE Std. 802.3.1-2011 <http://standards.ieee.org/about/get/>

<sup>5</sup> IEEE 802.3.1 incorporates MIB structures formerly defined in RFC 1213, RFC 2108, RFC 2863 and RFC 4836, relevant to IEEE 802.3 defined interfaces.

Finally, devices should report the maximum useful bandwidth as reported during the ATIS TEER bench test along with the tested power at the defined usage levels in the test. Each device may relate the real time cumulative throughput (sum of egress bandwidth across all interfaces) to the real time energy consumption in a product and vendor-specific manner.

CSCI strongly recommends that when possible all the above-described features be supported in newly procured enterprise routers and switches to enable network managers in performing their work in an energy-aware fashion.

## Key Power Management features for Attached Devices

In order for a network manager to act in a truly power-aware manner, network-attached devices must also provide power-related reporting mechanisms. While this is true to a degree for independently powered network-attached devices (e.g., NAS, servers, desktop computers), it is particularly important for Ethernet devices powered directly from the network using Power over Ethernet (PoE) as defined in IEEE 802.3. The energy footprint for these devices is directly within the energy usage of the network. Therefore the power monitoring and control for these devices should be managed in addition to the networking devices discussed above.

Even with the limited power provisioning & management capabilities of early-generation PoE equipment, the savvy network manager can engineer her/his network to reduce energy waste. The first generation of PoE (defined by IEEE 802.3af) includes rudimentary management and allows up to 13W of power to be supplied to the Powered Device (PD) from the Power Sourcing Equipment (PSE). This was updated by IEEE 802.3at to increase the power to 25W and add sophisticated management and control. The control mechanism allows the PSE and PD to negotiate power limits dynamically as needs and resources change. This ability is particularly important for PoE systems that allow a higher maximum power. If the PSE is designed with an assumption that all ports constantly supply the maximum power to PDs then the PSE power conversion will operate in a very inefficient region. Using the sophisticated control and negotiation of power modes, the PSE can optimize its power conversion for the particular power demands at any point in time. Furthermore, PSEs may benefit from more efficient design if there is no need to supply maximum power on all ports simultaneously (such a design is termed “oversubscribed”).

### **Attached Device Management example using Power over Ethernet (PoE):**

A PSE with 48 ports may offer to supply 30W of power on each port (i.e. supporting PDs up to 25W). If a fully subscribed design is used, the PoE alone would require 1.44kW from the power supply. The Power Supply Unit (PSU) of the PSE may be very efficient at (or near) full load, drawing only 1.6kW from the main supply. However, in practice, the PDs attached will not all require 25W simultaneously. For instance, a common installation may be supporting IP videophones with a maximum power close to 25W, but with an on-hook power requirement of less than 2W. When all the phones are on-hook (a normal situation) the total power draw is only

96W. This is less than 10% of the provisioned power and as typically designed will cause the power supply to run in a much less efficient part of its power delivery range. The power wasted by power supply inefficiency may be as much as 40% of the power used (i.e. 38W). If the power supply is provisioned to support only 25% of the phones drawing full power simultaneously, the maximum power capability is only just over 400W. Therefore the power supply will be running at 25% load when all of the phones are on-hook. This will reduce the wastage by more than half (i.e. 20W).

Much of the energy management for PDs may be controlled through the PSE, which should be supporting networked energy management as described above. Additionally, PDs should report their own energy and performance parameters in the same manner as defined for networking devices above. With real-time reporting enabled by both PD and PSE, a network manager has access to data that can form a foundation to guide network design, operations, and continual improvement.

Other energy using devices that are attached to the network may be monitored or managed through mechanisms appropriate to the device type. In some cases it may be useful to manage attached devices by proxy through the networking device management interface. The IETF Energy Management Working Group is also studying the measurement and control of energy for attached devices using a parent-and-child proxy mechanism. Particularly, this can be used to monitor, analyze, and coordinate the control of otherwise disparate systems (e.g. security systems and lighting) in order to maximize energy savings or to overcome problems of scalability for large numbers of controllable instances. In this sense, the network can be a fundamental enabler of broad energy monitoring and management in the enterprise, inclusive of ICT and other device categories.

### **Example for Coordinated Management of ICT and Other Equipment**

An office system has network connections for the security door entry system; the cubicle lighting systems; the cubicle wall power; and the desk telephone (an IP phone that is normally connected to the network). When an employee enters the building – as determined by the security system’s access control mechanisms - the management system is alerted, causing the lighting, power and phone at his workstation to be switched on. Similarly, when he leaves, the security system alerts the energy management application which powers down the equipment. This application has the advantage of saving energy without demanding that the employee changes his/her habits while also offering the convenience that the workstation is fully powered and operational by the time the employee reaches it. Such a system could save substantial power per cubicle for the entire time that the desk is unoccupied.

## Optimizing Network Energy with Application Requirements

Both the device level and the network level may have opportunities for energy-savings according to the network requirements of the applications in use on the network. Although devices may be expected to conserve energy as the network traffic load is reduced, such changes are generally dynamic and may be both brief in duration and unpredictable in the short term. Additional savings might be achieved if the device and network behavior can be optimized for energy efficiency according to the actual requirements of the applications currently active. While detailed definitions of such optimizations are out of the scope of this paper, one can identify some ways in which application behavior will reflect in the network's energy use.

### Network Load

The first type of optimization that may be coordinated using knowledge of the application requirements is based on foreknowledge of the expected load. At the device level, there may be subsystems that can be adjusted to reduce the power usage if the required maximum throughput is known to be lower. Similarly, at the network level, multi-path configurations that are required to support very high throughput can be reduced or disabled in order to save energy. There may be many reasons behind the variation of the required application load and the devices that are making the adjustment may know these reasons or an external entity that has oversight of the application requirements and the range of network adjustment may adjust the device behavior based on its available information.

### Examples

To illustrate this variety of optimization, consider the following examples:

(a) An aggregation switch (N) in the network that has a fixed number of 10 Gb/s Ethernet configured in an uplink fashion and numerous 1 Gb/s Ethernet links (M) connected downstream to leaf nodes such as PCs, servers, WLAN access points, and other switches. By looking at the current and historic activity level of its switch fabric including activity on downstream links, the aggregation device can compute the average aggregate bandwidth of traffic that is being switched upstream on the uplinks. If the load computed is less than the capacity of the upstream links, a subset of those links can be powered down to save energy without impacting network performance. In the event that there is a spike in network activity upstream, the uplink ports in the energy saving state can be brought back into an active traffic mode.

(b) An Ethernet blade server in a data center has two 10 Gb/s network interfaces equipped with offload capability. By looking at the anticipated tasks that are pending for the offload engine in conjunction with historic data, the server can place one of its interface links in an energy saving state and rely solely on the other networking interface until such time as there is a spike in the tasks that the server has to deal with.

A number of variations of the above theme can be conceived where a particular interface or set of interfaces is put into an energy savings state based on anticipated load. In addition other parameters, such as time of day which represent times where the load is known to be lower (e.g. an enterprise environment at night time), can be incorporated in the decision making process.

## Application Requirements and Power States of the Device

Another type of optimization depends on knowledge of the specific network requirements of the application in conjunction with deeper power saving states of the device. For example, certain applications may have specific requirements for network latency; the energy saving potential of systems that support Energy Efficient Ethernet may be improved if they are allowed to negotiate longer wake times; clearly, if the wake times are allowed to increase to the maximum that can be supported for the applications in use then the energy savings potential of EEE can be maximized.

### Example<sup>6</sup>

To illustrate this type of optimization consider the following example:

A wiring closet switch with a Gigabit Ethernet interface is connected to a desktop PC in an enterprise environment. During periods of low activity, such as nighttime, the switch and the PC may go into a Low-Power Idle (LPI) state. In the direction where the switch is transmitting to the PC, the switch, based on a control policy, places the link in LPI. Absent data link layer capability, the PC would simply power down its physical layer interfaces achieving a level of energy savings.

Considering that the application is a user application that may tolerate additional latency (vs. a low latency server application for example), the PC may wish to power down additional subsystems such as the PCI express interface between the network interface card (NIC) and the south bridge in the PC by putting it in a lower power state such as those defined by the PCI standard. Nevertheless, this may require additional wake-up time beyond that of the physical layer. The PC can negotiate with the switch, via the mechanism defined in EEE, additional wake-up time. If the switch has additional buffering resources, which is not uncommon for switching fabrics with pooled memory resources, it can allocate additional wake-up time to that port. When the switch transitions the link state from LPI to active, it sends the wake signal to the PC at which point both PHYs transition to the active state in the pre-determined physical layer time. However, the switch will hold off sending any packets to the PC until the full negotiated time has expired thus affording the PC additional time to transition its internal circuits, in this case the PCI express, back to the active state where it is ready to receive and process information. The net result is additional power savings beyond the physical layer.

There are a number of variations of the above where two link partners can take advantage of this mechanism. It is noteworthy that the mechanism is dynamic, allowing the link partners to renegotiate wake times based on a change in application, resources or other conditions. Furthermore, as the link is full duplex, the same mechanism applies in the other direction where the PC is transmitting to the switch.

---

<sup>6</sup> Refer to [http://www.ieee802.org/802\\_tutorials/2009-07/802%20july%20energy%208.pdf#Page=37](http://www.ieee802.org/802_tutorials/2009-07/802%20july%20energy%208.pdf#Page=37)

## Other Optimizations

Other optimizations may emerge with the coordination of application requirements and the variation of these requirements with the control of energy saving features of the network and its components. In all cases, the availability of control and monitoring of energy features will improve the development and prove the effectiveness of the optimizations.

## Detailed Management Requirements

The fundamental assumption in this section is that, Network Systems will consume full power when delivering expected work. Energy savings will come from stopping energy wastage during non-peak utilization. To determine the energy consumption, system behavior needs to be monitored. While the details below apply to a single Enterprise class network device, the concept scales to a network of devices.

### Variables to Monitor:

1. System Energy consumption
  - a. System power consumption at the inlet power (sampled with time stamp)
  - b. Power delivered by the power supplies (sampled with time stamp)
  - c. Power consumption of sub-systems or FRUs (sampled with time stamp)
  - d. Max Capacity of the Power Supplies
  - e. Power delivered per supply (if there are multiple power supplies in the system)
2. System Performance
  - a. Input data arriving at all system ports (sampled with time stamp)
  - b. Output data exiting from all system ports (sampled with time stamp)
3. System Environment
  - a. Inlet and outlet temperatures (sampled with time stamp)
  - b. Power consumption of Fan sub- system (sampled with time stamp)
  - c. Fan Speed (sampled with time stamp)
4. System Usage
  - a. System On/Off times
  - b. Start/End times for other power modes (stand-by, sleep)

Some of the applications of the above information are to understand the Device behavior with respect to energy consumed and the useful work done. This information would allow the user to see scenarios in which the device is operating at sub-optimal efficiency. Listed are a few examples where this information can be used.

1. The inherent nature of design makes certain Network devices consume energy at a constant level, regardless of throughput. The devices should be designed such that the energy consumption is proportionate to throughput.
2. Redundancy in power supply scheme along with low utilization; causes power supplies to operate inefficiently, hence wasting lot energy in power conversion. The power supply subsystem should be designed to operate at high throughput levels. Modular scheme should allow for adding or removing power capacity based on need. Thus keeping power supply utilization at high levels, which allows power system to operate at high efficiency.
3. Fan subsystem consumes substantial amount of power. If left unchecked, power will be wasted when system does not need high airflow for cooling. Fan speed control should be built into the system and fan power should be minimized at low temperatures.
4. Unused and unconnected portions of a device consume same amount of power as operational portions. A mechanism to shut down unused ports/subsystems will save a lot of energy.
5. Systems stay powered on when not in use. There are many use cases where devices are operational for a specific period of time. Automated networked applications can be employed to shut down unused devices.

The MIB module should have the necessary items in order to monitor these parameters. Most information pertaining to the system performance is already available. This information has been described in the Appendix A.

Note: Network and server equipment often use power supplies with a PMBus<sup>7</sup> interface. Newer PMBus specifications define an optional energy counter field so that the power supply itself can maintain energy consumption counters. Where those data are available, they should be incorporated into monitoring/reporting by the Device vendor.

Calculating Energy Efficiency: Using the measures of useful work and the measure of energy consumed. The network equipment should be able to report the energy efficiency of the system. This information can also be relayed to higher-level management software to make system level decisions. Alternatively, fundamental data that go into determining efficiency can be reported to higher-level management software that can then compute efficiency for individual devices, collections of devices, and the managed network as a whole.

---

<sup>7</sup> <http://pmbus.org/specs.html>

## Conclusion

The increase of networking traffic, combined with growing energy prices, is leading to greater emphasis on reducing energy use in networking devices. This objective is shared by environmentally conscious businesses and entities. This paper demonstrates how the design of energy efficient systems is a collective responsibility of system designers, developers, network infrastructure managers and other individuals who are involved in the design and operation of network systems. The focus of this white paper is power management and control, which is one of the key elements in achieving this objective.

Emphasis has been placed on real time monitoring of the energy, which provides a better understanding usage patterns and identifying potential areas of improvement. The capability to monitor several key parameters is already available in the current systems, while a few additional useful parameters have been identified. This paper categorizes the devices into two areas: “The networking devices” and “The attached devices”. The requirements of these two categories have been discussed. Finally, application level optimizations have also been specified. A few practical examples have been presented to demonstrate the potential value add that these policies offer.

Since network and infrastructure managers play a pivotal role in implementing the energy efficiency features, they should adopt the best known methods (BKM) to make their designs more energy efficient. Additionally they should communicate their requirements and provide feedback to the designers of networking equipment to improve the next generation designs. The system designers & developers should understand the industry direction and place emphasis on incorporating capabilities that lead to more intelligent, power-manageable devices.

This publication has been developed by Climate Savers Computing Initiative, the international member-driven consortium working to reduce the energy consumption of ICT (information and communication technologies) by increasing the energy efficiency of computing equipment, increasing the adoption of power management, and providing education about the environmental and financial benefits of energy efficient computing. Since its launch in 2007 the non-profit organization has grown to almost 700 members. Nearly 11,000 people have joined as individuals by taking the CSCI pledge to use power management and to purchase energy efficient computing products. The organization offers an online catalog of PCs, laptops, servers, power supplies and power management software that meet CSCI technical specifications for energy efficiency at <http://bit.ly/ba9PDy>. CSCI is led by Cisco, Emerson Network Power, F5 Networks, Google, Intel, Juniper Networks, Microsoft, Samsung and World Wildlife Fund.

We recommend the two other CSCI publications in this series on “Energy Efficiency” of Networking Equipment, on the topics of power supply and power conversion (PS BKM) and network-level energy savings (Network BKM).

## References

- U.S. Environmental Protection Agency, Energy Star Program. 2007. *Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431*. August 2.  
[http://www.energystar.gov/ia/partners/prod\\_development/downloads/EPA\\_Datacenter\\_Report\\_Congress\\_Final1.pdf](http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf)
- Koomey, Jonathan. 2011. *Growth in Data center electricity use 2005 to 2010*. Oakland, CA: Analytics Press. August 1. <http://www.analyticspress.com/datacenters.html>
- Lanzisera, Steven. 2010. *Data Network Equipment Energy Use and Savings Potential in Buildings*. September 1, 2010. ACEEE Summer Study on Energy Efficiency in Buildings. Lawrence Berkley National Laboratory. LBNL-3510E  
<http://escholarship.org/uc/item/2998x42q#page-14>
- IEEE Standard for Information Technology, IEEE Standard 802.3az*, 2010.
- IEEE Standard for Information Technology, IEEE Standard 802.3*, 2008.
- The Internet Engineering Task Force (IETF), IETF Energy Management Working Group  
<http://datatracker.ietf.org/wg/eman/>
- IEEE Standard for Management Information Base (MIB)*, IEEE Standard 802.3.1, 2011.
- IEEE 802 Tutorial, *Energy Efficiency and Regulation*. 2009. July 13.  
[http://www.ieee802.org/802\\_tutorials/2009-07/802%20july%20energy%208.pdf#Page=37](http://www.ieee802.org/802_tutorials/2009-07/802%20july%20energy%208.pdf#Page=37)
- PMBus Specifications. <http://pmbus.org/specs.html>

## APPENDIX A

The following list indicate the SNMP MIBs complying with RFC 1213 and RFC 2863 that provide performance indication.

- System up time: sysUpTime "The time (in hundredths of a second) since the network management portion of the system was last re-initialized."
- Bytes (octets) received: ifInOctets "The total number of octets received on the interface, including framing characters"
- Bytes (octets) transmitted: ifOutOctets "The total number of octets transmitted out of the interface, including framing characters."
- Unicast packets transmitted: ifOutUcastPkts "The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent."
- Unicast packets received: ifInUcastPkts "The number of subnetwork-unicast packets delivered to a higher-layer protocol."
- Non-Unicast packets transmitted: ifOutNUcastPkts "The total number of packets that higher-level protocols requested be transmitted to a non-unicast (i.e., a subnetwork-broadcast or subnetwork-multicast) address, including those that were discarded or not sent."
- Non-Unicast packets received: ifInNUcastPkts "The number of non-unicast (i.e., subnetwork-broadcast or subnetwork-multicast) packets delivered to a higher-layer protocol."
- Multicast packets transmitted ifOutMulticastPkts "The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses"
- Multicast packets received: ifInMulticastPkts "The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a multicast address at this sub-layer. For a MAC layer protocol, this includes both Group and Functional addresses."
- Discarded transmit packets: ifOutDiscards "The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space"
- Discarded receive packets: ifInDiscards "The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space"

- Transmit packet errors: ifOutErrors "The number of outbound packets that could not be transmitted because of errors."
- Receive packet errors: ifInErrors "The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol."

Additional information that is required that is not currently available on all systems:

- Timestamp for performance measurement
- Peak Performance rating for the device
- Number of samples read (for power Measurement)
- Peak Power Rating per Module
- Accumulated energy (i.e., Watt-samples)
- Number of energy accumulator overflows
- Instant power consumption (Watts)
- Chassis inlet air temperature
- Chassis exhaust temperature
- System on and off timestamps