climate
savers
computing SM

Energy Efficiency Guide for Networking Devices

Table of contents

Executive Summary

This document describes methods for comparing the energy efficiency of networking platforms. In particular it covers the comparison of networking devices that share a similar functionality or position in the network (a device class) using standardized metrics. We also aim to demonstrate how network designers and procurement personnel can assess the efficiency of systems from different vendors as part of the product evaluation and purchase process and make intelligent use of energy-specific product features.

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 though 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 (Koomey, 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 white paper is the final entry in the series and focuses on evaluating and comparing the energy efficiency of network devices for purposes of both procurement and network design. The information conveyed here is intended to be usable in both pre-production and field environments and appeals to network designers, operational managers and procurement personnel.

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
William Wael Diab,	Broadcom Corporation
Kathleen Fiehrer	Intel
George O. Goodman	Climate Savers Computing Initiative
Daniel Kharitonov	Juniper Networks

The authors wish to thank Pierre Delforge, Natural Resources Defense Council, Center for Energy Efficiency Standards, for his review and feedback.

Layered Model for Network Energy Savings

According to ITU-T recommendation L.1400 and FG-FN publication "Overview of Energy savings of Networks", network-related energy savings can be functionally separated into *applications* (services), *network*, *device* (element) and *component* layers (Figure 1). At the very top, applications and services are provisioned over existing communications infrastructure. Although high-level ICT applications rarely have a direct energy footprint, they are generally considered to have strong potential to reduce or increase indirect energy use in areas of communications, transportation, material handling and office productivity [GeSI 2008].

One level down, communication networks are material assets, uniting network devices into interconnected webs (core, access, corporate LAN, data center etc). Physically, networks may belong to different organizations such as public offices, private companies, content service providers (CSP), internet service providers (ISP) and others. It is quite common for one organization to maintain several networks of different types. It has to be noted, that although each network physically consists of network devices, its energy efficiency depends not only on constituents (elements), but also on the way they are interconnected (architecture) and on the features provided (functionality, redundancy, etc). Almost all significant network properties may affect its final energy outline, which makes this layer an independent target for optimization. The full profile of network energy use may also include the energy impact of attached devices – such as IP phones, alarms, video cameras and smart sensors.

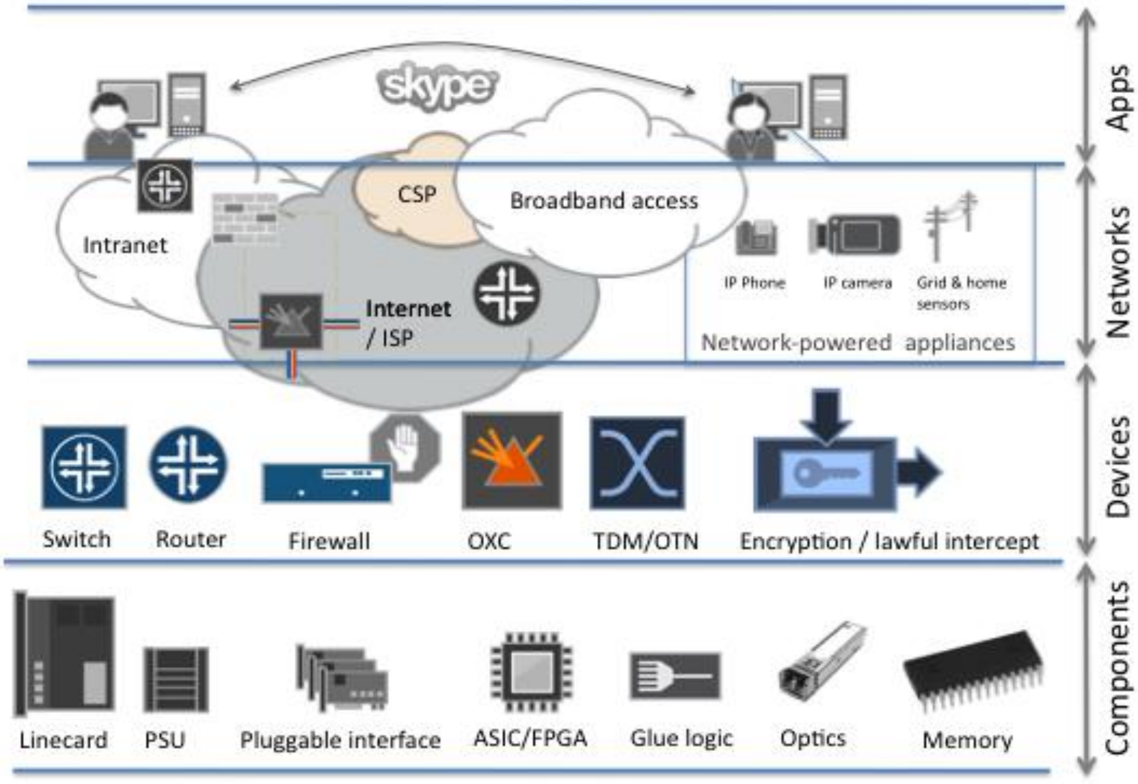

Figure 1. Layered model for Energy Savings. Source: CSCI

In the *device* layer, role-specific devices form the “parts bin” for designers and solution architects to devise networks from. Unlike networks themselves (which are mostly turn-key solutions fit to unique business requirements), network devices are strongly clustered into distinct categories (such as core, edge routers, Ethernet switches, firewalls, deep packet inspection systems and many others). Subject to strong competition from vendors, this layer is exceptionally suitable for optimization during the procurement (vendor selection) phase and is also the main target for organizational energy policies (runtime power management).

Finally, the bottom layer of energy optimization consists of *components* underpinning the network devices. This layer includes all sorts of merchant and proprietary silicon and technology from integrated circuits to lasers, power supplies, mechanical parts and fans. Again, within each component category, parts may compete on price, performance, size and other properties (including energy efficiency).

However, unlike members of other layers (which operate under end-user control and supervision), components are not tradable in the consumer market. In other words, end-users typically receive components deeply integrated within network devices and rarely (if ever) can choose between competing component designs.

This brings us to an important conclusion: the minimum (atomic) unit of energy efficiency exists at a network device (system) as a whole; energy-aware consumer decisions to adopt and use specific network-based solutions can be exercised only from the network device layer and up. This document focuses on energy efficiency of network devices; information on upper layers (such as energy-efficient network architectures) will be available in a future series of CSCI white papers.

Definition of Energy Efficiency

The top-level definition of energy efficiency is the amount of work delivered per quantum of energy. With the main function of network devices being transporting and processing of data (in bits) using energy supplied through the electric grid (in Joules), the measure of energy efficiency for network devices is naturally expressed in bits/Joule or Joules/bit[1]:

Efficiency = Energy / Throughput

where “Throughput” denotes the amount of data successfully transported by a network device (cumulative egress bandwidth). Note, that this definition is generic, and various network devices may employ significantly deviant definitions of “throughput” based on the class and role (e.g., a device performing analysis or processing of data in transit vs. one providing relatively simple switching). For example, reference throughput of a firewall needs to be computed under different conditions than that of a cache server or a packet switch.

---

[1] Including derivative denominations, such as W/Gbps, Gbps/kW and others

It goes without saying that the ultimate measure of efficiency would be the ability to use least energy for the same work in the real deployment conditions. However, reliable statistics from a live network may come too late (or be too expensive to collect) to inform an organization's purchasing decisions.

This is why, for purpose of product comparison, it is important to assess energy efficiency as early as possible, preferably starting from proposal evaluation phase. Such assessment can be done using pre-measured estimates – otherwise known as "energy metrics". To understand the value and applicability of such metrics, we first need to understand the use of energy within network devices.

Energy Use in Network Devices

Network devices perform a number of operations that can broadly be classified in network-specific or generic computing terms. Both of these categories of operations consume energy. As the scale of network applications continues to increase, power consumption of the underlying network devices tends to become a more pronounced portion of the power consumed throughout the end-to-end digital infrastructure.

Although end-users most often do not have control over embedded components, advanced knowledge of energy utilization can be helpful to gauge vendor's claims on the overall system efficiency. In order to provide the reader with representative examples, the typical energy use breakdown is outlined below.

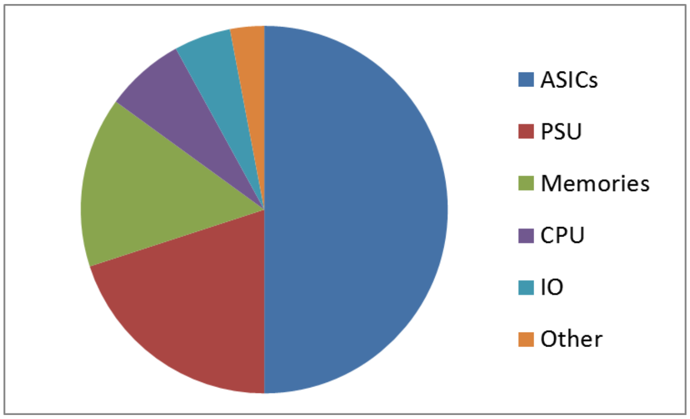

Figure 2. Energy allocation within a typical backbone router. Source: CSCI member companies.

In the example shown in Fig. 2, Application-Specific Integrated Circuits (ASICs) account for the biggest portion of the energy with about half of the power consumption being attributed to them. Power conversion inefficiencies account for about 20% of the overall energy usage. Memory could be in the 10-20% range depending on the device's configuration. The CPU and I/O components each contribute between 5-10% to power usage. In this particular network device, improvements in ASIC fabrication technology or degree of integration would give the most gain to the overall system efficiency, with memory architecture (for example, using RAM versus TCAM for packet lookup) being a second optimization target. On the other hand, power subsystems in this class of devices tend to be extensive but fairly effective with modest potential for significant improvements.

Next, we can look at the example of a device with a slightly different energy usage profile (Fig. 3).

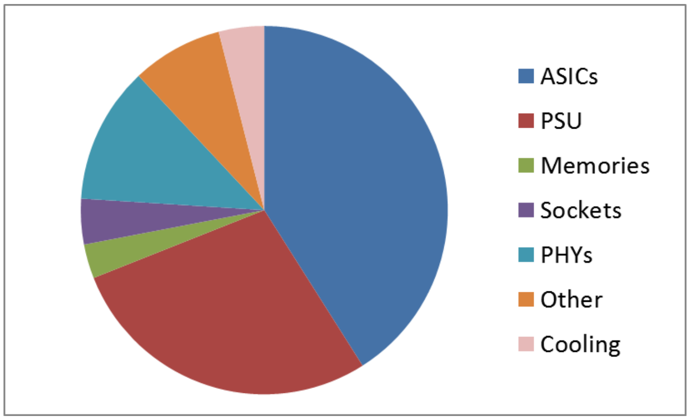

Figure 3. Energy consumption in a typical edge switch. Source: CSCI member companies.

Here, again, ASICs are responsible for the lion's share of energy, but power supplies now account for a greater energy percentage (25% to 30% range) and, due to device class, are more likely to be over provisioned. This makes the power subsystem more tempting to optimize (see [CSCI PSU 2011]). In this particular product category, physical layer interfaces (PHYs) also consume a noticeable amount of power (about 10%), which makes this device a good target for PHY-specific technology such as Energy-Efficient Ethernet. On the other hand, memory and cooling parts each consume fewer than 5% of energy and are not an optimization priority from a system energy use perspective.

Therefore, during the procurement cycle, it is helpful to remember that the best decisions should result in the lowest overall cost of ownership and a network device marketed for efficiency of a single subsystem (such as variable-speed fan tray or a modular power supply) may not be the most economical system overall (when factoring other components in).

Cumulative improvements in the subsystems and components detailed above are better reflected in system-level metrics, which aim to accurately assess the overall efficiency of a network device.

System-Level Efficiency Metrics

Energy reporting practices for network equipment have evolved with the concerns and priorities of network equipment users. Most commercial systems today are required to display their maximum power usage in a product datasheet. This is often referred to as "nameplate power" or "power rating". The nameplate power is normally used for power distribution sizing and generally remains much higher than the typical configuration of the device would ever approach. Besides, power rating tells nothing about maximum system throughput. Therefore, efficiency comparisons made on the basis of nameplate power can be misleading. In case of modular equipment, vendors may offer "power configurators" – tools that can be used to calculate (add together) the worst-case power consumption for systems with custom hardware configurations. The numbers generated by these tools are principally similar to nameplate power and are generally not at all useful for energy efficiency judgments.

Since both throughput (e.g. in bits per second) and energy consumption (Joules/sec) of telecom devices can be practically measured at the same time, this forms the basis for system-level efficiency testing and representation of results. As we have noted in the previous sections, such tests should not be burdened with internal details of architecture other than type and purpose of the device (i.e., such tests should use a black-box approach).

Inasmuch as the test profile is concerned, the optimal response of a network system should demonstrate energy consumption proportional to load. Quite similar to ideal computing devices [Barrosso 2007], ideal network equipment should consume negligible power when idle and only draw up to the power rating at peak load times. Although many network devices today show little change in their energy consumption with respect to load (and hence can be adequately described using *peak* efficiency metrics), newer energy-proportional technologies (such as Energy Efficient Ethernet) are encouraging the development of network systems with an improved dynamic range of energy-use and the use of more sophisticated *variable-load* efficiency metrics (such as ATIS TEER [ATIS 2009] and ECR-VL [ECR 3.0.1]).

Finally, network environments where temporary trade-offs between energy draw and performance can be acceptable, may find *extended-idle* metrics useful to approximate energy response with non- real time energy management. A typical enterprise may recover a significant percent of energy routinely consumed by the network devices during off-peak hours; however, this requires an energy policy designed to minimize the risks of service degradation from applying pre-scheduled "deep sleep" states for selected devices.

We can conclude this section with a notion that the constellation of three system-level metric types (*peak, variable-load* and *extended-idle*) describing various system-level energy responses together forms a complete basis for product comparisons. However, an attentive reader must already ask two questions: what system configuration the metric(s) should correspond to, and how the standardized load profile is relevant to real-life energy performance.

The first question mostly stems from the fact that modern network devices are often modular and can be equipped with "pluggable" modules of different utility and purpose. Quite obviously, the number of possible combinations of removable modules can quickly exceed the cost-information space required to collect and report them. Therefore, vendors are reasonably expected to measure and report the most energy efficient configuration of a given product as *certified*, with efficiency of other device configurations implicitly inherited or explicitly *declared* (e.g. by adding together module-specific energy footprints[2]).

The second question is based on the observation that no two networks are alike: it is pretty much guaranteed that the short-term (instantaneous) and the long-term (day/week) traffic dynamics will invariably differ between installations. This is why any efficiency metric built around a reference utilization profile (such as US EPA mile-per-gallon rating for automobiles [EPA fuel 2011]) is not guaranteed to match real-life energy performance of a given system under test: the latter can be higher or lower based on the use conditions. The key here is that a solidly designed metric should adequately represent system's baseline technology so that a device with a better value on the metric(s) is likely to consume less energy.

Introduction to Test Methodologies

State-of-the-art efficiency metrics for network devices use direct measurement of power drawn under class-specific load conditions. Such conditions vary according to the metric design: in some measurement cycles the network equipment is required to be able to transmit data at the highest rate at any time; in others the network load is designed to resemble extended low-utilization periods (such as nights and weekends). Additionally, the maximum offered load itself is a subject for determination because different network systems may exhibit vastly different overall performance.

Measuring energy consumption in relation to system throughput should be sufficient to compare efficiency of systems that perform identical functions; comparisons between systems with dissimilar purpose obviously make less sense. It also goes without saying that the quality of a certified metric will depend on the quality of test methodology, and it is the latter that allows results to remain robust and reproducible across different vendors and test laboratories. While every metric comes with its own test methodology, hereby we briefly review the concepts of energy testing.

Recommended Test Topologies

As previously mentioned, the first step in energy efficiency testing is the probing of maximum system performance. During such probing, the offered load (traffic) can be provisioned in multiple ways. A fully meshed traffic pattern is characteristic for a device that is positioned to operate in a network core. In this configuration traffic entering any port is forwarded to other ports in order to distribute the average load evenly amongst the ports (Figure 4).

---

[2] *declared* ratings may remain limited in scope and accuracy (e.g. provide only peak efficiency data)

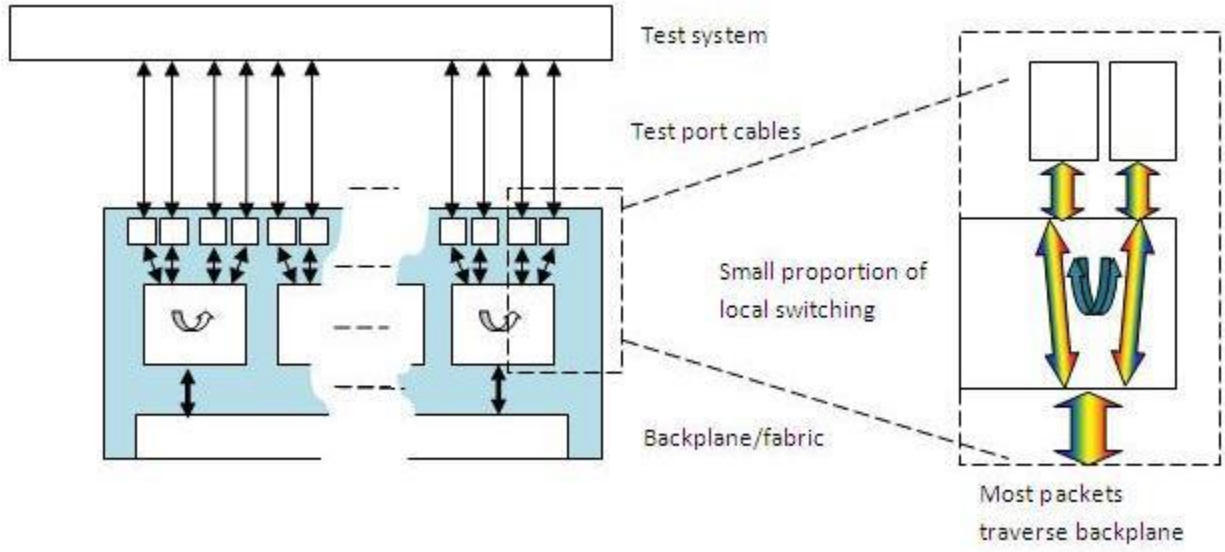

Figure 4. Full-mesh (core) test topology

On the contrary, devices that are operating towards the edge of the network do not tend to see such a large proportion of traffic between “client-side” ports, but rather forward packets from the downlink ports to the uplink ports and vice-versa. For this reason, most metrics also define an edge-specific traffic topology (Figure 5) that mimics the typical use in the access or aggregation network.

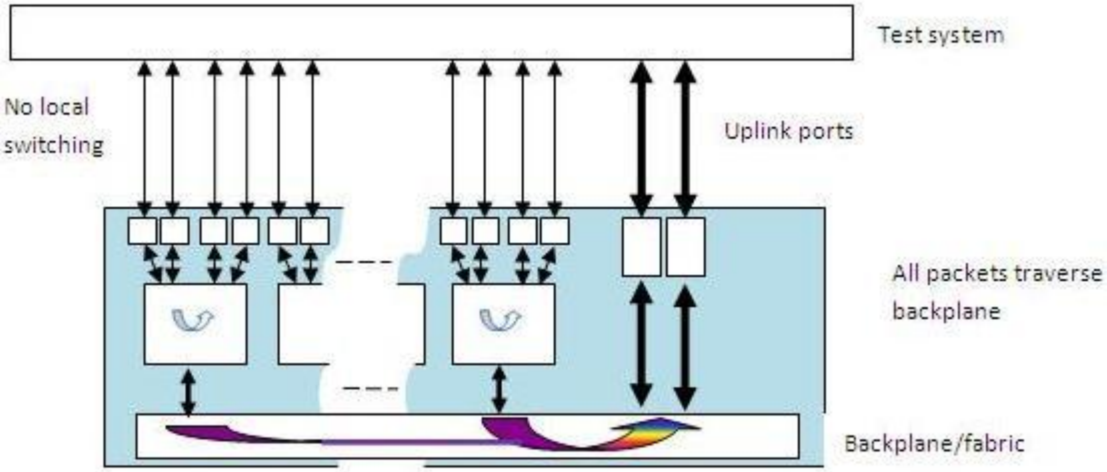

Figure 5. Partial-mesh (edge and aggregation) test topology

In the edge (partial-mesh) configuration the packets arriving on the uplink ports are destined for all of the downlink (edge-facing) ports with a random distribution so that the load average on each edge port is similar; the packets received on the edge ports are destined for the uplink ports in a similar way to ensure equal average load amongst the uplink ports.

Non-Recommended Test Topologies

A vast variety of test topologies can be conceived to constrain the effort and cost of passing traffic through system under test at proper scale. Though this is not an exhaustive list, such topologies may include static power measurements without passing traffic, traffic loops within the device with external or internal loopbacks ("snake test", "TTL multiplier test", and similar assessments with improper use of traffic generators/analyzers), mesh topologies not coherent with device class and usage (e.g. tests maximizing local switching capabilities on aggregation devices) and so on.

Metrics generated by such tests tend to lack the minimally required documentation and references to national or international standards bodies (such as ITU-T L-series recommendations) and thus should not be considered as valid data sources. The same consideration applies to any energy efficiency data provided without an explicit reference to any well-known and peer-reviewed test methodology.

Traffic Configuration

In addition to the forwarding patterns, there is some variation in the definitions for content of the traffic load. Firstly, the forwarding paths can be defined using either layer-2 or layer-3 addressing (Ethernet MAC, IP addresses, MPLS LSPs or similar). Typically this will be governed by the nature of the devices being tested – switches will use MAC addresses to control their forwarding paths, routers will use routing databases and so on. The test may require that address or route learning is performed prior to the start of the test. Additionally, the test may make use of layer-2 or layer-3 services such as virtual local area networks (VLANs) or quality of service (QoS) settings for the traffic being processed.

Some networking devices also have some form of "deep processing" function which may require certain amount of input traffic to be redirected and treated in a way significantly different from plain forwarding. If the test requires such processing, then the offered load will need to be constructed accordingly. Such requirements are rare in efficiency metrics specifications at this time but should be noted for future considerations.

The type of data used for the test payload can also vary in a number of ways. The test may challenge the system under test with specific packet sizes or may define a reference traffic mix. Typically, such mix will be picked to represent the traffic that would be expected in real applications. If the test does not define the packet size then it should be expected that the system developers would choose the packet size or packet mix that is most flattering for their particular systems. In some cases, the test may also define packet encapsulation, addressing schemes and/or fill patterns to fully exercise the forwarding path. Moreover, deep-packet inspection and stateful functions such as Crypto VPN, firewall and NAT/PAT functions may need complex, flow-driven combinations of layer 4-7 payloads.

System Efficiency Calculations and Proxy Metrics

A set of measurements collected in a controlled test environment typically produces an array of samples. On the other hand, a metric is typically a scalar, which requires using some formula to resolve constituent parameters into one number.

The classic definition of efficiency should match work done to energy used; however there is no easy definition of work done in a networking device outside of a specific role (or class). This latter fact sometimes brings the matter of "proxy metrics" – ways to assess energy efficiency via indirect performance indicators.

One frequently considered proxy metric is "set target" (power limit) according to which functions are supported (e.g. multiband wireless network, nominal port count, etc.). Clearly this method does not account for differences between the effectiveness of various devices and thus it would only be appropriate for simplest applications (such as consumer-grade networking without performance targets, see [EU CoC 2008]). Another proxy metric example would be the measure of some internal component utilization – such as the clock rate of system processor. A third proxy metric example would be judging the efficiency of an entire network device by the efficiency of a single subsystem or component (e.g. a power supply).

In general, all proxy metrics violate the "black box" approach to energy testing: some devices may have identical "nominal" set of features while practically being of different scale and quality. Besides, suitable proxy characteristics can be hard to locate or justify (e.g. for many modern network devices, CPU load is not a suitable efficiency proxy because traffic can be forwarded entirely in packet-specific silicon bypassing CPUs). Therefore, proxy metrics are not a good substitute to system-level efficiency testing and must be avoided where possible.

This is why today the majority of quality work estimates for network systems are expressed in terms of data bits or packets (or work rate in terms of bits/second or packets/second) successfully passed through the system (a.k.a., "egress data rate"). The conditions of input data and the required actions (header lookup, packet encapsulation, deep-packet inspection etc.) should be clearly defined in any test methodology. A calculation of efficiency should take into account only the minimum-required overhead.

For example, let's assume the system under test is an Ethernet aggregation platform with twenty-four Gigabit Ethernet ports (LAN PHY) on access side (towards Ethernet CPEs) and two 10Gbps Ethernet ports on the uplink side (towards MPLS core network). If this system can forward the incoming L2 frames (256 bytes each) towards MPLS core with egress interface utilization of 100 percent while drawing 350 watts, its peak efficiency becomes 350/19.922 = 8.8 Watts/Gbps. Variable-load efficiency (as seen over a mix of load profiles) will depend on system elasticity – but will very likely be lower (see Table 1). Meanwhile, an extended-idle metric (with unused components disabled) is quite probable to yield better results compared to real time variable-load metric, but at a cost of being usable only in pre-planned traffic conditions.

Measurement	Peak Metric	Variable-load	Extended-idle	Non-standard
100% (19.92Gbps)	350W	350W	350W	350W
50% (9.96Gbps)	-	290W	200W	-
25% (4.98Gbps)	-	220W	170W	-
Metric value	8.8 W/Gbps	17.2 W/Gbps	12 W/Gbps	7.9W/Gbps

Table 1. Sample metric lineup

It may also happen that a vendor chooses to report a non-standard efficiency metric obtained with a proprietary or undisclosed methodology. In the example above (last column of Table 1), we illustrate an attempt to synthesize a non-standard efficiency metric by dividing known peak energy consumption by arithmetic sum of face port values. Such an approach may result in higher nominal throughput (relative to value measured with a recommended test topology on Fig. 5) and thus can be utilized for unfair marketing advantage. This is why attempts to declare efficiency information lacking public and peer-reviewed methodology should be clearly rejected.

Complementary Efficiency Features

A network device may have complementary energy-related features that do not affect its own footprint and hence are not reflected in system-wide efficiency metrics.

One example is the ability to report runtime energy utilization and load: while indispensable for energy planning, assessment and advanced power management, this feature (in itself) is merely an enabler for policy applications. For another example, an explicit energy-aware control plane services may allow a network to dynamically balance its multi-path resources according to traffic requirements in such a way that energy is saved when application load is reduced. However, an individual device that takes part in this optimization may not benefit directly from the energy savings[3].

More examples can be drawn from a notion that a network system is part of a larger web of appliances and may coordinate and enable energy savings in connected devices.

A prime example of a feature that saves energy on the local device but also enables energy savings on an attached device is Energy Efficient Ethernet (EEE) (IEEE 802.3az). The energy-saving functions of EEE are only enabled when the local device and its link partner both support EEE (although the link can still communicate if only one partner supports EEE). When the link partner uses its EEE functions on transmit paths, the local device can then save power on its own receive paths. Such savings will not be reflected in network device’s own energy metric because they occur on connected devices working in consortium. This relation goes both ways - the degree of energy savings possible for the local device will depend on the aggressiveness of the EEE function on the link partner. Similarly, the local device may choose a more or less aggressive policy for its transmit path that will dictate the level of possible energy savings for its link partner.

---

[3] Although energy-aware routing and load balancing technologies such as [Yong 2009] are currently not considered mainstream, they may become viable in the future based on balance of cost, risk and savings

Networking devices may also support network proxy functions (such as those defined in ECMA-393 or IEEE 802.11v). These functions allow network-connected devices to go into a sleep state and stay asleep while the proxy responds to routine network interactions (such as ARP or neighbor discovery). The ability to support network proxy again does not improve the networking device's energy efficiency (as measured) but nevertheless allows the attached appliances to save energy by sleeping more deeply or for longer periods.

Yet another feature that enables link partner savings is the power management feature built into Power over Ethernet (PoE) (IEEE 802.3at). This feature allows two PoE devices to negotiate power levels according to application requirements and resource constraints. In the absence of PoE management, the power sourcing equipment (PSE) is forced to assume that the powered device (PD) may use any amount of power up to its maximum rating at any time. This forces the PSE to over-provision the power reserved for PoE, resulting in an unreasonably large power supply that operates in a very inefficient manner. The PoE management feature, when supported, allows the PSE to manage its power supplies in an oversubscribed configuration. It also allows the PSE to control the PD power usage during a power outage, reducing the need for over-provisioned and wasteful UPS resources to guarantee basic service coverage.

Presentation of Energy Properties

As part of our mission, CSCI works to promote industry progress in the energy efficiency of network devices. We do this by offering a standardized disclosure program to the vendors of network devices. This disclosure is meant to be a way to communicate energy features in a robust, easy-to-use format, similar to information required on some consumer products (e.g. automobiles or washers/dryers) by environmental regulation authorities. The fully provisioned CSCI disclosure includes four distinct sections, shown by example in Figure 5.

In the first section of the CSCI-certified disclosure, the vendor is required to clearly identify and print the model name (part number) and configuration of a certified system under test.

In a second section, the CSCI disclosure contains energy metrics obtained for the reference configuration. At present, the minimum requirement is disclosure of an ATIS TEER metric (variable-load), with optional peak (ECR/EER) and extended idle metric (ECR-EX) values[4]. Whenever a metric is disclosed, conditions (such as the device class of the system under test) and metric's constituent components (measured values of performance, consumption and utilization) should also be clearly displayed.

In a third section, complementary measurements and efficiency features are listed. At present, CSCI disclosure format requires listing two features: system's ability to report its own real-time energy consumption and support for IEEE 802.3az.

The minimum compliance level for the first feature is the ability to gauge runtime energy consumption (within 1 second or better of resolution) of a whole system at least at power distribution system and being able to communicate it via SNMP, command-line interface, XML or other means of remote monitoring. Being able to monitor power consumption right at the grid inlet and report per-component power distribution is also preferable (but not required for feature compliance).

---

[4] More than one metric value for each class can be displayed if device transcends a class definition

The minimum compliance level for IEEE 802.3az is to be able to support this feature on any number of fixed or modular (pluggable) EEE-compliant Ethernet ports that may be installed on a listed device (or chassis).

The support level for either of the two complementary efficiency features should be expressed as “supported”, “not supported” or “not applicable”.

In a fourth section, a vendor is allowed to use a limited amount of freeform information to list additional and proprietary energy-related features. This section is limited to three lines of plain text with a font size equal to or smaller than that of sections one, two and three. The fourth section should be positioned in the lower (bottom) half of the CSCI disclosure leaflet.

The entire disclosure should not contain any graphics, logos and text beyond the sections outlined above or supplied by CSCI. A sample disclosure leaflet is provided below (Fig 5).

The accuracy and validity of information in the disclosure remains the vendor’s responsibility; CSCI reserves the right to deny participation in the disclosure program to vendors violating the program terms and conditions[5].

ACME AC3572-EX, 72-port 10Gb Core Ethernet Switch

ATIS TEER = 313, ECR = 3.38 W/Gbps

Peak reference throughput = 719Gbps, draw @100% utilization = 2,433W; @30%=2,320W, @0%=2,180W

Inline energy monitoring	
yes	✓
no	
N/A	

IEEE 802.3az	
yes	✓
no	
N/A	

SmartEnergyManager™ compatible

Figure 5. Sample CSCI disclosure format*

* Please bear in mind that graphical rendition and layout of the actual CSCI disclosure is subject to change without notice. The latest information on this program is available from www.climatesaverscomputing.org

---

[5] Please refer to CSCI web site for complete terms of disclosure program

Practical Use of Energy Efficiency Data

The Climate Savers Computing Initiative disclosure program for network devices is intended to provide valuable information for procurement and network design/operation personnel, acting as an enabler for monetary, environmental and operational improvements and savings. In this section we briefly summarize the relevant considerations.

Using Metrics in the Procurement and Network Design Process

According to Dell'Oro, the average selling price of a fixed L2/L3 switch across the top five vendors in first quarter of 2011 was US$43 per port, which equates to US$1,034 per 24-port device configuration. Such a switch would typically consume between 50 and 75 watts of energy. Considering useful life for such a device of 5 years and average energy footprint of 60 Watts, this sample device would consume 2,630 kilowatt-hours in that 60 month period. Doubling this consumption to account for minimally needed cooling and AC/DC conversion capacity would increase this figure to 5,260KWhr. Assuming the typical U.S cost of electricity ($0.112 per kilowatt-hour per the Department of Energy's 2010 report) and factoring an estimated 5% annual increase in electricity bills, the electric grid-related cost of ownership to the end user will mount to $650 - over 60% of the initial equipment acquisition cost.

In other words, the cost of electricity needed to run a device as simple as desktop Ethernet switch over the course of five years can offset capital discounts and concessions originally offered by vendors. As an illustration, an estimated energy efficiency improvement of 20% for this equipment type would be equal to 13% in capital savings (a value of $135) – a number that can become significantly higher in countries with higher electricity rates.

With performance, features and cost of acquisition being the leading modalities of modern network devices, the procurement process typically starts with a list of products that offer comparable utility and value. Energy efficiency, being a second-order product characteristic, nevertheless should work as a tiebreaker, indicating a product's ability to conform to the end-user's operational expense (OpEx) expectations, energy footprint, and environmental policy. A comprehensive, standardized energy disclosure program (such as the one offered by CSCI) makes this process straightforward and allows for formal comparison of energy-related capabilities during RFP evaluation.

Metric Significance in Relation to Deployment Conditions

However, there might also be cases, where the conditions of deployment may prompt the customer to look for information beyond the minimum required in CSCI disclosure. Below is a brief list of conditions that may need additional investigation or information to be provided by the vendor:

- Configuration notably different from that of reference (disclosed) system. As we discussed in the section "System-Level Efficiency Metrics", this may warrant request for additional information (including certified or declared metrics) to account for difference in custom configurations. This situation also arises when equipment under consideration uses many diverse linecards (hardware modules) and/or is configured below 25% of capacity (see Fig. 6 for example of fill impact on system metric validity).

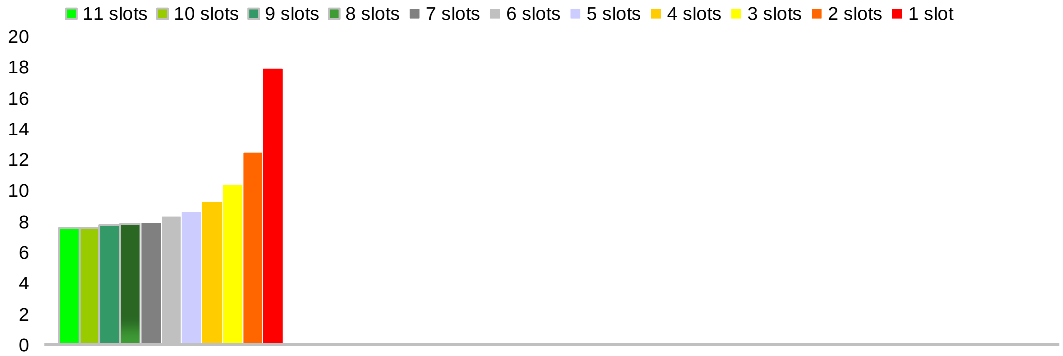

Figure 6. Sample peak efficiency (W/Gbps) with respect to chassis slot fill. Source: [TDAEE 2009]

- Substantial, measurable energy reductions enabled beyond the network device itself. In section "Complementary Efficiency Features" we have provided some examples of client-side energy savings that can be enabled by the network device. In the environment, where such indirect gains can be prevalent over the network footprint, the relative weight of system-level energy efficiency metrics in RFP decision-making can be proportionally reduced

- Single device performing significantly divergent network functions. Whenever a network device can transcend class definitions and be redeployed or co-deployed for an additional role, its efficiency may change significantly (see Fig 7). In that case, peer-reviewed test methodologies recommend publishing two (or more) relevant metrics. The final efficiency of a hybrid device (for instance, a router that also acts as a firewall) can be higher or lower than that of stacked single-purpose devices, depending on those metrics and the percentage of traffic affected.

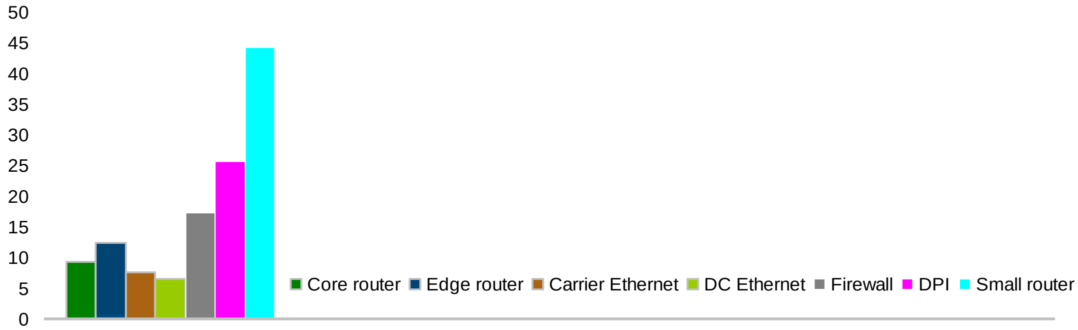

Figure 7. Sample peak efficiency (W/Gbps) for devices in different classes. Source: [TDAEE 2009]

Disclosed Metrics in Relation to Runtime Energy Performance

In operational (field) environment, network equipment can be monitored *in situ* for runtime efficiency using external probes or embedded utilization and energy monitors [CSCI PM 2011]. Whenever such monitored network system comes along with CSCI disclosure, it is possible to record its energy performance and verify against the reference data points.

In general, the observed efficiency should remain comparable to metric values disclosed for matching load (throughput) points. Larger discrepancies may signify suboptimal equipment hardware inventory (e.g. poorly utilizing a modular chassis) or software configurations – the latter especially important when verifying extended-idle metrics on a device with active idle energy management policies. However, in the runtime environment it is also important to honor the consequences of system utilization levels, which may be actually different from those reference levels published in CSCI disclosure. As evident from Figure 8, network efficiency drops exponentially with decreasing load, which makes it hard to observe the expected efficiency levels whenever equipment is underutilized[6].

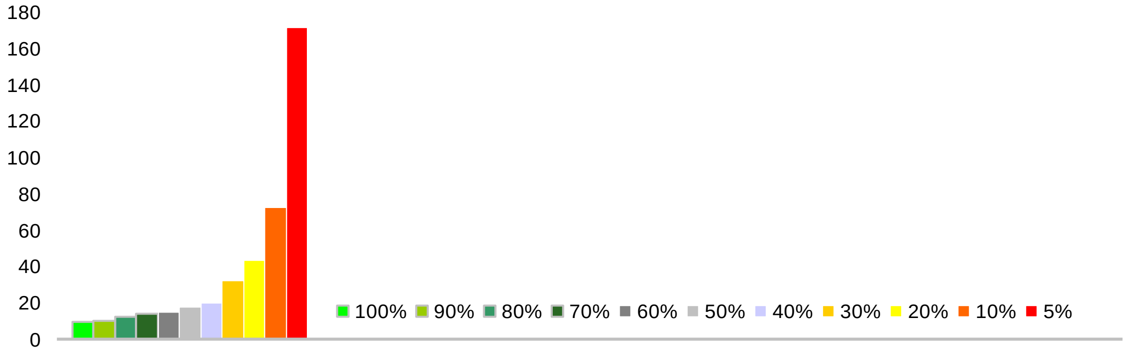

Figure 8. Peak efficiency (W/Gbps) as function of system utilization. Source: [TDAEE 2009]

Nevertheless, observing and quantifying runtime efficiency can offer very valuable insights into how the network is using energy and potentially uncover ways to improve efficiency of existing installations by consolidating traffic, removing unused equipment and running the energy policies. In other words, CSCI disclosure helps to establish a baseline for energy performance and improve field operations up to and beyond the reference measure.

---

[6] When calculating runtime efficiency and comparing it against the reference profile, one may also need to pay attention to other factors – such as reference methods for calculating throughput and bursty nature of traffic in live networks (where instant utilization may have high deviation parameters)

Conclusion

In this white paper, we discuss practical aspects of energy efficiency in network devices and highlight energy savings potential both inside and outside the network fabric. We demonstrate, that energy efficiency of network devices is suitable for optimization based on solid and quantifiable arguments in both procurement (vendor selection) and operational (energy management) phases. We present an original CSCI disclosure program for network equipment specifically designed to arm customers with reliable, no-nonsense efficiency information furnished by responsible equipment vendors. Finally, we discuss the role of standardized efficiency information in RFP evaluation and provide guidelines for operational assessment of energy efficiency.

References

[ATIS 2009] “Energy Efficiency for Telecommunications Equipment: Methodology for Measurement and Reporting for Router and Ethernet Switch Products” ATIS, standard ATIS-0600015.03.2009; Mar 2009

[Barroso 2007] Barrosso, LA and Holzle, Urs; “The Case for Energy-Proportional Computing”; Computer; Dec 2007

[CSCI PM 2011] “Power Management for Networking Devices” Climate Savers Computing; Oct 2011

[CSCI PSU 2011] “Considerations for Selecting Power Supplies for Networking Equipment and Evaluating Power Conversion Efficiency” Climate Savers Computing; Oct 2011

[ECR 3.0.1] “Network and Telecom Equipment—Energy and Performance Assessment. Metrics, Test Procedure and Measurement Methodology” ECR Initiative; 2010

[EPA fuel 2011] US EPA fuel economy test method page http://www.epa.gov/fueleconomy/regulations.htm#test

[EU CoC 2008] “Code of Conduct on Energy Consumption of Broadband Equipment” EUROPEAN COMMISSION; Nov 2008

[GESI 2008] “Smart 2020: Enabling the low carbon economy in the information age” Global e-Sustainability Initiative; 2008

[Koomey 2011] Jonathan G. Koomey "GROWTH IN DATA CENTER ELECTRICITY USE 2005 TO 2010," Analytics Press report, Aug. 2011

[LBNL 2011] Nordman, Bruce and Lanzisera, Steven; “Electronics, Networks and Energy” IEEE ICCE; 2011

[TDAEE 2009] D. Kharitonov “Time-domain approach to energy efficiency in high-performance network element design” IEEE Globecom; 2009

[US EPA 2007] Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431. Prepared for the U.S. Environmental Protection Agency, ENERGY STAR Program, by Lawrence Berkeley National Laboratory. LBNL-363E. August 2, 2007.

[Yong 2009] Yong Wu, Chiaraviglio, L, Mellia, M and Neri, F; “Power-Aware Routing and Wavelength Assignment in optical networks” ECOC; 2009