

Energy Efficiency for Network Equipment:

Two Steps Beyond Greenwashing



August 10 2008



Table of Contents

Executive Summary	3
Introduction.....	3
The Basics of Moving Data	4
Step One: Efficiency Criteria.....	4
What to Expect from an ECR Metric	5
Decoding Vendor Datasheets.....	5
Advanced Topics in Efficiency Criteria.....	7
Metric Accuracy and Practical Impact.....	8
Step Two: Design Goals	8
Conclusion	9
About ECR Initiative	10
Literature.....	10

Executive Summary

Rising energy costs and increasingly rigid environmental standards work in tandem to draw attention to the “power efficiency” aspect of data networking. Governments and corporations around the world are tightening energy and emission budgets, thus creating demand for new, energy-aware generations of telecom equipment.

In response, telecom vendors are starting to label their offerings as “green” and “environment-friendly.”

But one important detail about this theme is often missing: verifiable data to support these “green” marketing claims. This paper aims to fill the resulting gap by discussing the practical aspects of achieving and objectively measuring energy efficiency in the telecom world.

As the first step, we define an efficiency metric that supports informed decisions related to network equipment power consumption. In the remainder of the document, new criteria are used to define the best energy practices and the practical aspects of energy efficiency for the networking industry.

Introduction

Daily price increases at the gas pump are typically the world’s most recognized symbol of rising energy costs. Expensive fossils also mean expensive electricity, which is highly relevant in the corporate world, where operating costs are rising commensurate with the local utility bills.

Meanwhile, most people believe that reducing energy-related office and transportation expenses is fairly straightforward, driven primarily by media and government agency popularization. Advanced lighting technology, heat management, and low-impact transportation systems are expected to lower the cost of everyday living—albeit possibly at the expense of a higher initial expenditure. In fact, the latter frequently determines the right point to invest in energy-saving technologies.

In order to discuss energy efficiency in telecom networks, let’s imagine a day in the life of a Chief Green Officer (CGO) for a large corporation.

To date, all employees have received energy-efficient laptops, and overhead daylight panels have been replaced with low-power organic LED matrices. The entire building has been fitted with an advanced airflow system that substitutes for the legacy forced-air cooling. Furthermore, a pilot project has just begun to integrate solar panels into windows and glass walls, the first step towards a zero-sum energy operation.

One morning, our CGO drives her hybrid car to work for a routine meeting with the Chief Technology Officer (CTO) to discuss an upcoming corporate network expansion. But the first thing she hears upon arrival is very startling: an expensive network overhaul will likely increase the company’s energy footprint by several hundreds of kilowatts—a disheartening message in direct contrast with the improvements the company has already implemented.

Anxious to achieve her energy conservation goals, she calls one leading telecom vendor for comments. What she hears goes along the lines of: “Measuring the consumption of each network component is not very useful, you should look at the overall solution to determine the final efficiency.”

Promptly failing to construct the “overall solution” in her head, she tries yet another vendor and gets lucky this time. Rather than being brushed off, she gets an interesting perspective on how an assorted collection of power management technologies will help her network become energy-efficient in the not-so-distant future. But yet again, she fails to obtain firm numbers or commitments.

Sound familiar? This is the current state-of-the-art of power management in the networking industry.

Unlike the cost of transportation and light bulb ownership, energy consumption for telecom equipment is not a simple metric; it depends on many parameters, most prominently technology, performance, and applications. This mix of parameters makes it challenging to estimate the actual energy efficiency of network equipment—a situation clearly undesirable for customers, government agencies, and companies focused on innovation.

Solving this problem requires a coordinated effort of vendors, governments, and customers alike to identify and clarify metrics that unambiguously and objectively define the energy efficiency of the network world.

The Basics of Moving Data

It is widely accepted that modern data storage and transmission were made possible by the discipline of informational theory, established by Claude E. Shannon in his landmark article "A Mathematical Theory of Communication" in 1948. As a practical consequence of his work, an arbitrary message can be represented via units of informational entropy. Modern computing and communication devices typically use binary entropy units (bits) to store and access information within the device, although serial communication links can represent data in more complex forms. In any case, information exchange happens by alternating and reading the unit states, a process that requires electronic or optical gates to transition between different energy levels.

Power consumption in networking equipment is related to loss during the transfer of electric charges, which in turn is caused by imperfect conductors and electrical isolators. The exact rate of this consumption depends on technology (operating voltage and fabrication process), as well as the frequency of transitions and the number of gates involved. The latter is driven by the architecture and purpose of a telecom or network device.

However, the number of circuits, gates, and their individual operating frequencies are not very interesting to end users.

Instead, when choosing equipment for a new network build-out, the list of requirements for the data plane is typically comprised of a set of features (minimum packet processing requirements), scaling parameters (bandwidth, queues or other revenue-generating assets), packaging format and relative cost. Vendors are free to offer any platforms and technologies that satisfy the respective requirements. Thus, a typical Request for Proposal (RFP) simply defines a product class (for example, a core router or an Ethernet switch), and equipment manufacturers offer products of variable capacity that fit this profile.

Step One: Efficiency Criteria

As different vendors develop competing technologies and architectures, this leads to unequal power consumption between equipment belonging to the same class. Therefore, in theory, to determine the winner in energy consumption, it would be sufficient to put two or more network devices under the same load and measure their respective power draw.

In reality, however, this is rarely possible.

First of all, full-scale testing of network equipment requires a non-trivial investment in test gear, possibly costing millions of dollars in the case of high-end routers and switches. Second, network devices come in different capacities, and this raises a question as to what should be the system configuration for the test. Fully configured platforms of unequal scale will obviously produce energy readings that are not directly comparable.

In order to define an efficiency metric, we need to normalize energy consumption E by effective full-duplex throughput T . This will give us a normalized energy consumption rating,

$$ECR = \frac{E}{T}$$

where E denotes the energy consumption (in watts) and T denotes the effective system throughput (in bits per second) [SAINT 2008].

The values of E and T may come from either internal testing or the vendor's data; in both cases, they should be verifiable.

For the current generation of network equipment, *ECR* is most conveniently expressed in watts/Gbps, identifying the amount of energy (in joules) required to move an array of data (in bits) across the device. *ECR* can also act as an approximation of system consumption based on the number of active gigabit (or $N \times \text{Gbps}$) ports¹.

What to Expect from an ECR Metric

As we mentioned before, *ECR* can be a valid differentiator within a product class, where several vendors may contend for energy efficiency with competing technologies and system architectures. Equipment with better (lower) *ECR* ratings will spend less energy to move the same amount of payload.

This is where network equipment loses the analogy to light bulbs and passenger cars—the amount of commercial payload (system capacity) must be factored into efficiency estimates. If capacity requirements are growing, they may or may not be satisfied within the energy budget originally reserved for the legacy network. Internet traffic can, in fact, be compared to commercial freight—it naturally takes more energy to transport an increasing volume of cargo.

Contingent on standardization [ECR 1.0.2], knowledge of an efficiency metric immediately positions a product relative to its competition on the energy grid; this is informative and promotes competition.

But the value of a normalized efficiency rating goes well beyond head-to-head platform comparisons.

A standardized way of measuring energy efficiency paves the way for forward-looking requirements and goals. A typical network equipment vendor cannot comply with the goals set by a single testing laboratory—this does not scale, as every customer may come with a unique set of tests and network conditions. The picture changes with *ECR*, which represents a unified way of testing and can be used for defining long-term R&D targets.

On the other hand, standardization of *ECR* measurements will mean that each specific (customer) setting may not be met exactly. For instance, if a customer only plans to deploy 600 Gbps of switching capacity, and the considered platform is rated at 1.6 Tbps in a full configuration, the standardized *ECR* value will not match the actual efficiency numbers observed in the field. The same restriction applies to a standardized test load relative to any particular packet traffic.

However, there is a good chance that relative *ECR* standings between the platforms will stay the same across a wide range of configurations and offered load profiles—an important assumption that represents a compromise between custom and standardized testing. While the latter should be freely and widely available, the former may require significant investments, but yield better precision. We expect that for most practical purposes, a standardized *ECR* rating will be an adequate energy performance estimate.

Decoding Vendor Datasheets

The vendor-agnostic nature of the *ECR* metric makes it a good basis for head-to-head comparisons, but with one catch—in order to produce comparable metrics, parameters *E* and *T* should be collected with the same methodology and acceptable precision.

In theory, even if the network equipment vendor does not report *ECR* or a similar energy efficiency metric obtained under controlled methodology, it should still be possible to approximate *ECR* using the data publicly available in the datasheets or test results.

Let's first look into platform throughput.

The easiest thing to define, throughput equals the amount of data the platform can process per second under reasonable conditions; that is, traffic pattern and system configuration should be commensurate with the product class. Most often, however, vendors report platform capacity, not throughput. Capacity is typically given as half-duplex bandwidth, expressed in Gbps. This number, when divided by two, should normally give the full duplex throughput (*T*) of the system.

¹ Sometimes *ECR* can be also quantified in Gbps/Watt, which is an inverse of the same metric. In this case, it is referenced to as $EER = 1/(ECR)$
PUSHKAREVA-10-08-2008

Unfortunately, this is not always the case.

Some vendors express system capacity as the theoretical peak utilization the system can achieve based on a single isolated parameter (for example, switch fabric design). This may impose significant limitations—things like line cards that can't use their (potentially available) fabric bandwidth, inter-slot restrictions preventing the system from reaching the stated goal in a full-mesh traffic scenario, or some other limits with sustained load profiles and usability. Thus, vendor data should always be verified.

For the purposes of an *ECR* calculation, throughput *T* may only include effective (revenue-generating) capabilities; therefore, it should be calculated as a sum of the capacity of all active line cards in full-mesh configuration, while the traffic profile should correspond to the class of the device. For example, a stated 720 Gbps throughput of an edge router may have to be reduced to 240 Gbps or even less, if the platform under test is actually performing as such. Sometimes this type of data is hard to obtain and requires unbiased study of internal and external test reports.

Energy consumption is the part that is even harder to figure out from a vendor's collateral.

A typical datasheet or hardware guide may include several energy metrics, including but not limited to:

- Power System Rating (also power supply rating, expressed in amps or watts)
- Maximum Power Consumption (expressed in watts)
- Component-based Consumption Estimate (expressed in watts)
- Typical (Average) Power Draw (expressed in watts)

Power System Rating reflects the site preparation requirements recommended by the vendor and tends to be very conservative. Also known as "Power Supply Rating," it can potentially be used as an estimate for consumption, but with a possible error margin. In some cases, vendors outfit their platforms with high-capacity power supplies in planning for future system upgrades; the actual system consumption may be a fraction of what the power supplies can deliver.

Maximum Power Consumption can also serve as an upper boundary estimate for the power draw. However, it tends to penalize modular systems designed for the optional high-power components. For example, an Ethernet switch designed with Power over Ethernet (PoE) modules in mind will have much higher maximum power draw than a fixed copper-port model; yet both systems may yield identical consumption in a pure 1000 BaseT mode of operation. In addition, a maximum power consumption estimate can change without notice when new modules are introduced and old modules withdrawn from production.

Component-based Consumption Estimate allows a customer to estimate a power draw in a "customized" configuration by adding the configured parts (components) with known power ratings together. This can be a fairly precise estimate, assuming the vendor has published accurate power numbers for all base, optional, and physical interface modules in the current board revisions. The availability and usability of such data is vendor-dependent and requires a thorough knowledge of system structure and operation.

Typical (Average) Power Draw is often reported to offset the (otherwise unremarkable) numbers shown in the "Power System Rating" or "Maximum Power Consumption" section, and tends to gravitate to the conservative side of the power consumption range. Motivated to demonstrate the low current draw, vendors are free to report this metric with underpowered configurations, components or load profiles that yield the best results; omission of a published test methodology typically signifies these and similar issues. In the lack of a public disclosure on measurement conditions, "typical" or "average" power draw cannot be reasonably used to rate the device against any other platform.

In conclusion, we may say that if the vendor does not report the standards-driven *ECR* metric, reconstructing it with a platform's datasheets or hardware guides can be a challenging task. Independent test results may be required to obtain performance estimates in a controlled condition and configuration, and conservative "maximum power consumption" figures may have to be used in lieu of test-based measurements.

Advanced Topics in Efficiency Criteria

So far, we have focused on peak efficiency: *ECR* represents energy performance of a fully-configured system running at maximum load:

$$ECR = E_f / T_f$$

where throughput T_f corresponds to the effective (measured value) and E_f corresponds to energy consumption under the full load.

It is not unreasonable to question whether a peak metric is relevant to more typical applications, where a network platform may run at a fraction of the rated capacity.

There are two distinct situations here: the first is related to configured system capacity. Modular telecom platforms are rarely deployed in full configurations from the start; instead, they typically reach their service ceiling midlife, after the network has gone through expansion and upgrade rounds.

To estimate the effect of a partial configuration, we can represent the power draw of a modular router or switch to be a sum of a fixed part F (chassis, host system, fabric, clocking) and a variable part V (which represents removable line cards, interface ports, and physical line drivers).

$$E = F + V$$

It is trivial to demonstrate that a system with more efficient fixed and variable parts (as normalized by throughput) in a full configuration will also remain more efficient across all partial configurations. If this condition is not true, a crossover point can be found, where a previously less effective system may become more efficient with proportional reduction of removable components (typically, a fairly degraded configuration).

For most practical cases, partial configurations will never change the relative standing of comparable platforms. Moreover, a higher utilized system will yield better energy efficiency in the first place.

The more interesting scenario happens when a running system is not fully loaded with traffic; this case is very common across all network and topology types. Due to the inherently fractal nature of Internet communications, even the busiest networks in the world still have their on-peak and off-peak times.

The opportunity for hardware and software designers here is to exploit drops in packet rates to conserve the power consumed by the network device. Estimating this dynamic power capability would require measuring the power consumption under a variable load. Therefore, it would be useful to complement *ECR* with an optional synthetic (weighted) rating (*ECRW*) that reflects off-peak and idle conditions:

$$ECRW = ((\alpha \times E_f) + (\beta \times E_h) + (\gamma \times E_i)) / T_f$$

where E_f , E_h and E_i correspond to energy consumption in full-load, half-load and idle modes respectively. Throughput T_f corresponds to the effective value obtained during *ECR* calculation. Weight coefficients α , β , γ represent the relative importance of different modes of operation and are often given the following values: $\alpha = 0.35$, $\beta = 0.4$, $\gamma = 0.25$ [VZ.TPR.9205]

ECRW has no direct physical meaning and thus is dimensionless; from the practical perspective, if *ECRW* happens to be equal to *ECR*, this means the system under test has no dynamic power management capabilities. On the other hand, the system that ideally follows the Barroso's principle of energy proportionality [IEEE Computer 2007] should have been measured to $ECRW = (0.55 \times ECR)$. Of course, this ideal case is not achievable, yet *ECRW* can highlight the technological improvements the modern telecom systems will be bringing into the market.

Metric Accuracy and Practical Impact

Regardless of which energy-related metric can be constructed or reported, it has little merit by itself; it is merely a tool to improve efficiency in real-world network deployments. Just like any other tool, it also has the attribute of accuracy, which defines the standard divergence of results. In a physical sense, the accuracy of *ECR* and *ECRW* reflects the averaging done in the course of measurement and fluctuations within the device; the former relates to sampling frequency during the test run and the latter relates to technological deviations. Such deviations do not only mirror the tolerance of the components utilized by the vendor, but also cover minor revision changes the vendor may apply to boards and circuits.

The experiential evidence of energy consumption testing at Juniper Networks laboratories suggests that the practical *ECR/ECRW* accuracy typically stays within ± 2.5 percent, which accounts for metrology errors and platform variations. Doubling the margin for safety, we can conclude that a practical *ECR/ECRW* advantage can be reliably reported within 10 percent or higher difference.

The practical implications of *ECR* diversity can be expressed in two ways: the monetary difference and the environmental impact.

The monetary difference relates to the cost of ownership over the life span of the network equipment, which, in most cases, is between 60 and 120 months. If a carrier point of presence (POP) or mid-size data center consumes 100 kilowatts, an *ECR* difference of 10 percent is equal to 438,300 kilowatt-hours ($24 * 365.25 * 5$) over five years of ownership. This number reflects the raw consumption of the network equipment and does not incorporate collateral energy charges such as cooling and power conversion. For the purposes of a cost estimate, a fully burdened consumption should be considered, which typically ranges from 1.5 to 3 times the input of network equipment. Assuming relatively conservative 2x cooling and conversion overhead, starting kilowatt-hour cost of \$0.10, and a fairly optimistic case of 10 percent rate increase per year, our reference *ECR* improvement of 10 percent is equal to \$107,000 in reduced operating expenses.

Likewise, *ECR* difference of 20 percent would yield \$200K in operating expenses, 50 percent would equate to half a million in savings, and so on.

The environmental impact of network operations creates a second dimension for energy efficiency.

Following ratification of the Kyoto protocol, many countries have accepted CO₂ reduction targets relative to 1990 emission levels. As a result, energy-consuming industries have to comply with the new legislation. The monetary aspect of such a transition may have less impact compared to political values and administrative mandates that can cause some corporations to find themselves in dire need of energy-efficient technologies. Looking at immediate targets (2008-2010), *ECR* forms the basis for equipment selection criteria. When planning for future (2011-2020) goals, *ECR* becomes a set of values and common language to speak to equipment manufacturers, which can be critically important for success of environmental efforts worldwide.

In environmental terms, choosing equipment with a 10 percent *ECR* advantage can be equal to a one year advantage on the carbon emission roadmap; a 50 percent advantage can be as good as reaching a forward-looking target covering the lifetime of new network deployment. This is where technology intersects with ecology.

Step Two: Design Goals

Defining a set of energy efficiency criteria can be compared to putting a stake in the ground and drawing a set of coordinates around it. This first step is necessary to define the starting point and the direction of progress. Designing efficient network platforms is a second necessary step for the success of an energy-aware networking paradigm.

As a complex combination of software and hardware, a network platform of any kind cannot become energy-efficient by chance; instead, it has to be meticulously designed to achieve advanced functionality within a limited energy budget. While a technical discussion on energy-related design topics [EE Penguin] is outside the scope of this paper, the general process of technology improvements is worth a look.

At the highest level, energy-related improvements in network equipment design can be classified as organic and engineered. Organic efficiency improvements are commensurate with Dennard's scaling law—every new generation of network silicon packs more performance in a smaller energy budget.

Engineered improvements refer to active energy management including, but not limited to—idle state logic, gate count optimization, memory access algorithms, I/O buffer reduction, and so forth.

It is interesting to note that some passive (organic) and active (engineered) energy management enhancements go directly in step with building network systems, while some do not. For example, better density, integration, and heat management allows for building faster and denser platforms—a clear differentiator in the market. On the other hand, dynamic power management proportionate to an instant load is a technology that does not fundamentally affect platform density or capacity. However, such and similar intellectual property form a pool of promising technologies that may, in fact, improve energy efficiency at a pace exceeding Moore's law. The return on investment in such areas is not always material, while being a pioneer is a challenging and expensive role. Therefore, innovative companies clearly need support from corporations and government agencies for continuous research in the area of energy efficiency. We expect that industry success in this area will increasingly rely on the effectiveness of a dialog between researchers, engineers, and business leaders.

Another remarkable area that may bring a lot of difference is the discipline of network design, where customers, consultants, and value-added resellers can team for improved topology design, equipment selection, and expansion planning.

Here are some simple guidelines that can make a profound difference in the overall result:

- Select equipment with best ECR/ECRW ratings.
- Plan for good system and path utilization; avoid excessive capacity reserves.
- Include rack space, cooling, and power conversion in your assessment, as these items can make a big difference in monetary and environmental impact.
- Do not multiply entities beyond necessity. Universal “Swiss army knife” designs are often the worst in power draw. Lean and effective core, edge, and access infrastructure provide the best in energy efficiency.
- Consider local energy costs and trends—this can affect the total cost of ownership

To estimate the effect of such change, let's theorize what our CGO can do after getting in touch with a network vendor that understands the importance of energy management.

First, her company's new network project can be redesigned for a smaller energy footprint, lower OPEX and simplified structure—all factors strongly promising to offset short-term higher capital expenses over the long haul. Her forward-looking energy goals can also be re-balanced against enhanced network requirements in areas such as support for telecommuting, paperless operation, and telepresence. This can result in improved company workflow – without environmental compromise.

*One imaginary business case. One real planet.
Whether they can co-exist will depend on you.*

Conclusion

Various estimates of Internet traffic yield different predictions, yet everyone seems to agree that the importance of data communications and the volume of traffic will continue to grow. This growth will likely vary from a linear to an exponential curve, depending on the country and economic conditions. Still, we can reliably conclude that individual consumers and businesses will increasingly rely on data networks worldwide.

Does this mean that networking equipment will be responsible for a global increase in energy consumption? Not necessarily.

Indeed, modern networks are frequently growing faster than Moore's law, which is the main reason for the absolute increase in power consumption (also seen by our imaginary CGO in the first section of this paper). However, much of this growth is fueled by reduction in other areas like commuting, offline shopping, offline banking, and goods manufacturing. The social phenomenon of "connected life" is transforming our world for a better quality of living with less material footprint.

As we are moving towards an information-driven society, human activity increasingly shifts into online domain, and efforts to build sustainable and environmentally-friendly equipment will stay high on the priority list. Demand for noticeable and verifiable energy efficiency will only continue to rise.

The two much-needed steps are an industry effort to standardize telecom network efficiency, and the rise of sustainable network technologies.

For an increasing number of very valid reasons, it's clear that we need to take these steps.

About ECR Initiative

The ECR is an open initiative that welcomes participants and users from network equipment manufacturers, government agencies, carriers, and enterprises. It is immediately applicable to establish the relative energy efficiency of network components produced by different manufacturers within a solid, unified and vendor-neutral test methodology. The specification is free to the public and unlocks the potential of the minimum-resource network design, where the energy footprint is reduced in parallel with meeting and exceeding the operational objectives. Additional information can be found at www.ecrinitiative.org

Copyright ©2008 ECR Initiative. All rights reserved.

Literature

[VZ.TPR.9205] Verizon NEBS™: Energy Efficiency Requirements for Telecommunications Equipment. Todd Talbot, Ludwig C. Graff

[SAINT 2008] Proceedings of IEEE Symposium on Applications and the Internet:
Luc Ceuppens, Daniel Kharitonov, Alan Sardella, "POWER SAVING STRATEGIES AND TECHNOLOGIES IN NETWORK EQUIPMENT OPPORTUNITIES AND CHALLENGES, RISK AND REWARDS"

[ECR 1.0.2] Network and Telecom Equipment - Energy and Performance Assessment, Test Procedure and Measurement Methodology. Draft, version 1.0.2, ECR Initiative, 2008

[EE Penguin] Energy Efficiency - Penguin-friendly Networking in Design © Juniper Networks, Inc, 2008

[IEEE Computer 2007] The Case for Energy-Proportional Computing, Luiz André Barroso and Urs Hölzle