Network and Telecom Equipment - Energy and Performance Assessment

Test Procedure and Measurement Methodology

Draft 1.0.4, Nov 10 2008

Revision Record:

Version	Date	Author	Reason
0.2	12/12/2007	dkh	First draft
0.3	03/22/2008	dkh	Input from METI “Small Routers” document
0.5	06/11/2008	aalimian	Extended product classes, algorithm graph
0.6	06/15/2008	dkh	Changes per first ECR trials across different product types
0.8	06/28/2008	dkh	Input from Verizon VZ-TPR-9205 spec
0.9	07/15/2008	dkh	Changes and comments per input from Bruce Nordman (LBNL)
1.0	10/12/2008	dkh	Input from ITU-T Focus Group on ICT’s and Climate Change
1.0.3	11/06/2008	dkh	Clarifications in VPLS Carrier Ethernet class

Areg Alimian, IXIA Corp aalimian@ixiacom.com
Bruce Nordman, Lawrence Berkeley National Lab BNordman@lbl.gov
Daniel Kharitonov, Juniper Networks, Inc. dkh@juniper.net

Table of Contents

List of Acronyms

ECR - Energy Consumption Rating (scalar metric, Watts/Gbps)
ECRW - Energy Consumption Rating Weighted (synthetic metric)
ECR (class Y) - ECR metric computed for class equipment Y
T_f - Measured maximum effective throughput (full-duplex, Gbps)
E_f **-** Energy consumption under full load, Watts
E_h - Energy consumption under half load, Watts
E_i - Energy consumption under idle conditions, Watts
SUT - System under test

Purpose

The purpose of this document is to define a framework for first-order approximation of energy efficiency for packet-based network and telecom equipment.

Background

As energy efficiency in datacenter and network and telecom space becomes an topic of increasing interest, it spurs significant activity across commercial companies and government agencies to define the sets of energy efficiency criteria for qualification and homologation purposes.

So far such work has been mostly oriented towards private testing and qualification – an approach best suitable for early adopters and organizations with extensive qualification labs and test budgets. This document complements these efforts with an open efficiency criteria that can be independently estimated and reported by equipment vendors and test labs for ease of selection and comparison.

The purpose of this work is tri-fold:

1. Define a test procedure for measurement and estimate of energy efficiency for network and telecom equipment. This includes methodology and test classes.
2. Establish a common energy efficiency metric for the network and telecom industry
3. Promote energy awareness and competition between OEM vendors

Theoretical basis

This document defines energy efficiency as energy consumption normalized to effective throughput. Such approach goes in line with high-level methodology suggested in [SAINT 2008] [ATIS-0600015.xx.2008] [VZ.TPR.9205] documents.

In other words, this document assumes the more energy-efficient network system to be the one that can transport more data (bits) using the same energy budget (in Joules).

Note, that packet-based systems offer a specific challenge for this approach.

Since the amount of data the system can transport can be equal (or lower) than the theoretically possible limit, system performance should be subjected to measurement alongside with energy consumption.

Scope

This definition is best suited for medium- to large-scale network and telecom systems primarily serving data streams. It is less relevant to small office, SOHO and multipurpose devices, where throughput is less relevant and efficiency criteria need to be more complex and involved, such as described in [METI 2008] [EC CoC Broadband]

In general, this document is applicable to many types of network and telecom equipment, including, but not limited to – routers, L2/L3 switches, optical shelves, security devices, load balancers, etc.

Class-specific requirements and test profiles are detailed in Appendix B.

Measurement procedure

Network and telecom packet-based systems are fundamentally based on the notion of statistical multiplexing, where system performance may or may not correspond to the bandwidth theoretically possible based on the "face" port configuration. To take this into account, this test methodology purports to perform simultaneous performance and energy consumption measurements under the load profile and conditions typical to the environment, where the system under test (SUT) is intended to operate.
The details of the test suite and offered load are specific to product class described in Appendix B.

The actual measurement cycle is designed to be simple, fast and inexpensive to run. It can be fully automated and, whenever there is a room for interpretation, should be designed to reflect the utilization profile and conditions frequently experienced in the field.

There is no SUT configuration change allowed any time beyond preparation phase. All energy savings adjustments (if done) by the SUT should be automatic

The procedure consists of four major steps.

System under Test (SUT) Preparation.
SUT is configured according to class requirements and offered load defined in the class requirements (Appendix B). Prior to the actual test, the SUT has to be exposed to environmental conditions outlined in Appendix A for at least four hours to settle the potential temperature difference.

Router tester equipment is used to simulate the load and collect the performance-related results. AC or DC inline meters are used to calculate energy consumption during the test. Appendix A lists metrology-level requirements to DC and AC-based meters.

Step 1 (qualification)
First run determines the maximum load that can be sustained at zero packet loss. Any methodology is suitable, including binary search (similar to RFC2544), heuristics or known maximum load values. There is no time limit for this run.
The run is complete after a maximum (lossless) load is determined.

The following three runs should be separated with idle time of 300 seconds or less

If the test class requires the SUT to be "primed" with control plane information (ARP/MAC/route learning etc), it should be done within the idle time window.

Step 2 (full load)
Second run offers the load $\mathbf{L_{max}}$ (identified at step 1) to SUT for period of 1200 seconds[1]. Energy consumption is being sampled for the entire period, and average consumption $\mathbf{E_f}$ calculated[2].

Step 3 (half load)
Third run reduces the load $\mathbf{L_{max}}$ twice ($\mathbf{L_{half}}$ = 0.5 x $\mathbf{L_{max}}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period, and average consumption $\mathbf{E_h}$ calculated. Load reduction is achieved by reducing packet rate on all configured ports.

Packet loss during second or third run (if seen) invalidates the measurement and resets testing to first run to provide a better $\mathbf{L_{max}}$ estimate

Step 4 (idle load)
Idle run removes the load and runs for another 1200 seconds. Energy consumption is being measured for the entire period, and average consumption $\mathbf{E_i}$ calculated. Load reduction is achieved by idling packet rate on all configured ports, or disabling ports on packet tester side, at vendor discretion.

Metric computation

One primary and one secondary metric are calculated with results from the aforementioned measurement algorithm. There are two methods to convert load $\mathbf{L_{max}}$ into effective full-duplex throughput $\mathbf{T_f}$ (expressed in Gbps).

In first method, the measured $\mathbf{L_{max}}$ is decomposed into a packet-per-second rate and packet sizes corresponding to the load. If packet sizes are variable, the average proportions are to be computed. Next, all applicable **minimum** L2 and L1 overhead is added to compute the effective wire-rate, at which the SUT performed. Note, that the idle timeouts inserted by SUT to compensate for asymmetric test patterns are not counted for.

Example1.
The SUT is a packet platform that can drive ten 10GE ports at 7,291,702 frames per second each with 64B Ethernet frames without loss. According to the tester, this corresponds to 7,291,702 packet-per-second rate per each port.

---

[1] The measurement interval represents a compromise between accuracy and speed. Shorter time intervals (i.e. 300 seconds) may overestimate the system performance (via temporary packet buffering) or cooling capabilities (i.e. system tolerance to heat dissipation before the fans change speed) . Longer measurement intervals make the testing procedure more expensive.

[2] Please refer to Appendix A for measurement conditions and qualifications

T_f = *10 x 7,291,702 x 8 x (64 + 1 + 7 + 12) = 49.000237440 Gbps*
(accounting for Ethernet SOF, preamble and minimum IPG)

In second method, the tester equipment itself can report the highest achieved line utilization on per-port basis. In this method, the well-known line rates for selected transport interfaces are multiplied by port utilization to calculate the final data rate.

Example 2.
The SUT is a VPLS edge platform with 10x 10GE ports (LAN PHY) on access side (towards Ethernet CPEs) and 10x 10GE ports on the network side (towards MPLS core network).

The SUT can forward the incoming L2 frames (256 bytes each) towards MPLS core with egress interface utilization of 100 percent. However, because of the 1:1 matching of access and network sides, the access side can only be utilized at 99.22 percent to allow lossless application of the 2-byte MPLS L2 VPN header required for packets on the network side. Same limitation is seen in the opposite direction, where the incoming network-side packets can only fill the access-side interfaces at 99.22 percent after the headers are stripped.

Data rate for 10G Ethernet IEEE 802.3ae is 10,000 Mbit/s

T_f = *10 x 10.000 x 1 + 10 x 10.000 x 0.9922 = 19.922 Gbps*

Example 3.

The SUT is an Ethernet switch that can operate at 100 percent line utilization when configured for 802.1q packet encapsulation (VLAN headers applied). The same switch can only operate at sub-line rate speed when not configured for VLAN encapsulation.

The measurement results from the second case should be used, unless the test profile specifically requires VLAN encapsulation to be present.

The primary metric is a peak ECR value, which is calculated according to the formula:

ECR = $\mathbf{E_f/T_f}$ (expressed in Watts per Gbps)

Where: $\mathbf{T_f}$ = maximum throughput (Gbps) achieved in the measurement
$\mathbf{E_f}$ = energy consumption (Watts) measured during running test “Step 2”.

ECR[3] is normalized to Watts/Gbps and has a physical meaning of energy consumption to move one Gigabit worth of line-level data per second. This reflects the best possible platform performance for a fully equipped system within a chosen application and relates to the commonly used interface speed[4]

---

[3] Sometimes, energy efficiency is also reported in Gbps/Watt under the name of EER; EER = 1 / (ECR)
[4] Other denominations could also be used, such as Watts/10Gbps, Watts/100Gbps or Joules/Gigabit

Second metric is a weighted (synthetic) metric that takes idle mode into account. It is used in addition to the primary metric to estimate power management capabilities of the device.

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

Where: Tf = maximum throughput (Gbps) achieved in the measurement
Ef = energy consumption (Watts) measured during running test “Step 2”.
Eh = energy consumption (Watts) measured during test “Step 3”.
Ei = energy consumption (Watts) measured during test “Step 4”.
α, β, γ = weight coefficients to reflect the mixed mode of operation[5]

ECRW reflects the dynamic power management capabilities of the device, which matches energy consumption to the actual work accomplished. An ideal system following the Barroso’s principle of energy-proportional computing [IEEE Computer 2007] should be able to achieve ECRW rating identical to 0.55 * ECR. A system with marginal power management capability is expected to demonstrate ECRW rating to be very close to ECR.

Reporting

Results can be variably reported based upon a class definition, or a combination of application and packet size, such as: ECR (class A) = Y, or ECR (Class A, B) = Z, Where A = equipment class, B= payload type, x = packet size; Y, Z = calculated efficiency

For instance, ECR (Class 1.1) = 12 Watts/Gbps; ECR (Class 1.2, IPv6) = 5.2 Watts/Gbps

For comparison purposes, the data can be collected in tables to reflect head-to-head competitive situation typical to RFP qualification, for example:

	Product A	Product B	**Product C**	Product D
Product class	Core	Core	Core	Core
Nominal Capacity	640G	1.28T	1.6T	3.2T
ECR (C1.1)	15 W/Gbps	12 W/Gbps	**9 W/Gbps**	12 W/Gbps
ECRW (C1.1)	9.1	11	7	10

In addition to the final result, the full SUT configuration, software version and hardware board revisions should be provided, along with traffic generator/measurement tool passports and environmental conditions. The test setup should be fully described, including the choice of offered load and test conditions chosen at the discretion of the vendor within a range of possible choices (i.e. Class 1.1, IPv4).

Reported results without the required documentation are assumed to be invalid.

---

[5] In this document, we specify $\alpha = 0.35$,$\beta = 0.4$, $\gamma = 0.25$ [VZ.TPR.9205]

REFERENCES

[VZ.TPR.9205] Verizon NEBS™: Energy Efficiency Requirements for Telecommunications Equipment. Todd Talbot, Ludwig C. Graff
http://www.verizonnebs.com/TPRs/VZ-TPR-9205.pdf

[METI 2008] METI Final Report (Small Routers, L2 Switches) by Router, etc. Evaluation Standard Subcommittee, Energy Efficiency Standards Subcommittee of the Advisory Committee for Natural Resources and Energy
http://www.eccj.or.jp/top_runner/pdf/tr_small_routers-apr_2008.pdf

[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"

[EC CoC Broadband], Code of Conduct on Energy Consumption of Broadband Equipment, Draft Version 3 JOINT RESEARCH CENTRE Institute for the Environment and Sustainability
Renewable Energies Unit European Commission

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

[**NIPP-TEE-2008-031R4**] ATIS Standard on Transport Energy Efficiency
ENERGY EFFICIENCY FOR TELECOMMUNICATION EQUIPMENT:
METHODOLOGY FOR MEASUREMENT AND REPORTING
TRANSPORT REQUIREMENTS

APPENDIX A. Measurement Conditions

A.1 Temperature

The equipment shall be evaluated at an ambient temperature of 25°C ±3ºC. The SUT itself should stay offline or operate at this air temperature for no less than three hours prior to the test. No ambient temperature changes are allowed until the test is complete.

A.2 Humidity

The equipment shall be evaluated at a relative humidity of 30% to 75%

A.3 Air Pressure

The equipment shall be evaluated at site pressure between 860 to 1060 mbar

A.4 DC Voltage

The input to the SUT (all active feeds) shall be at a nominal DC voltage of ±5%

A.5 AC Voltage

The input to the SUT (all active feeds) shall be the specified voltage ±1% and the specified frequency ±1%

A.6 Metrology requirements

Every active power feed should have the power (amp) meter installed in-line, with desired accuracy no less than ±1% of the actual power consumption

A.7 Sampling frequency

Ef, Eh and Ei calculations are based on averaging multiple readings over the course of measurements. Power meter(s) should be able to produce no less than 100 evenly-spaced readings in every 1,200 sec test cycle.

APPENDIX B. Proposed Product Classes & Test Applications

Disclaimer: For the purposes of public testing, all platforms should be tested with publicly available (shipping) software images, publicly available (shipping) board hardware revisions and fully documented and supported configurations

Class 1 - Routers

C1.1 Core routers

Description. Core routing platforms are systems with Terabit (half-duplex) or higher capacity. They are designed to provide line-rate performance in network cores with minimum functions (packet lookup and forwarding/switching). Core routing platforms come in various form factors, in standalone and multichassis enclosures.
Qualification. 500G or better full-duplex capacity, IP/MPLS FRR or equivalent
Test Application: IPv4, IPv6 or MPLS forwarding at discretion of the vendor;
L3 packet size (MPLS considered L3): 46B; forwarding over any types of forwarding entries (static, connected, IGP, EGP) – no less than 64 active routes.
Interface types: SONET, 10GE or 100GE as designated by the vendor, SR optics
Redundancy. For the purposes of testing, all redundant components (fabric, routing engines, power supplies, memory cards etc) **should be present** in the system

Note. For purposes of C1.1 testing, the symmetric full-mesh topology (same-bandwidth traffic stream from every port to every port) is required to be configured as the offered load.

C1.2 Edge Routers

Description. Edge routing platforms
Qualification. MPLS VPN capability
Test Applications: IPv4, IPv6 VPN, PWE, or VPLS forwarding at discretion of the vendor;
Payload packet size: 238B (IPv4, IPv6 or Ethernet frames as delivered to/from access side are considered to be payload); forwarding over any types of forwarding entries across all VPN instances, no less than 2K VPN destinations active (PWE circuits, VPLS hosts, IP VPN routes)
Interface types: at vendor discretion
Redundancy. For the purposes of testing, all redundant components (fabric, routing engines, power supplies, memory cards etc) **should be present** in the system

Note. For purposes of MPLS VPN forwarding (VPLS, PWE, IPv4/IPv6 VPN), the ports on SUT can be grouped into "network" and "access" side ports according to vendor discretion. Every "network" side port should be configured to send to every "access" side port and vice versa

C1.3 Multipurpose routers

Description. Routing platforms of variable purposes (enterprise, edge, etc)
Qualification. L3 forwarding
Test Applications: IPv4 or IPv6 forwarding at vendor discretion. L3 packet size: 494B; forwarding over any types of forwarding entries, no less than 16K active routes.
Interface types: electrical or optical at vendor discretion
Redundancy. For the purposes of testing, redundant components *may be* not present

Class 2 - WAN/Broadband Aggregation Device

C2.1 BRAS devices

Description. Legacy broadband aggregation devices
Qualification. PPPoE, PPPoA, PPP termination, per-subscriber QoS
Test Applications: PPPoE, PPPoA, PPP forwarding at discretion of the vendor;
L3 packet size: 256B; forwarding over any types of per-subscriber entries, no less than 64K subscribers with no less than four (4) queues assigned to each.
Interface types: SR optical at vendor discretion
Redundancy. For the purposes of testing, all redundant components (fabric, routing engines, power supplies, memory cards etc) **should be present** in the system

C2.2 BSR/Common Edge devices

Description. Broadband aggregation devices, Ethernet-oriented
Qualification. PPPoE, PPP, IP DHCP, per-subscriber QoS
Test Applications: IP/DHCP, PPPoE, PPP forwarding at discretion of the vendor;
L3 packet size: 238B (measured as IPv4 or IPv6 payload to/from access side); forwarding over any types of per-subscriber entries, no less than 64K subscribers with no less than four (4) queues assigned to each.
Interface types: SR optical at vendor discretion
Redundancy. For the purposes of testing, all redundant components (fabric, routing engines, power supplies, memory cards etc) **should be** present in the system

Class 3 - Ethernet L2/L3 Switches

C3.1 Carrier Ethernet Platforms

Description. Carrier-grade Ethernet switching platforms
Qualification. MPLS L2/L3 VPN capability over Ethernet.

Test Application: IPv4, IPv6 VPN, PWE, or VPLS forwarding at discretion of the vendor;

Payload packet size: 238B (IPv4, IPv6 or Ethernet frames as delivered to/from access side are considered to be payload); forwarding over no less than 128 VPN instances in a full-mesh configuration between "network" and "access" port groups.

Interface types: SR optical (10/100/1000/10GE) at vendor discretion
Redundancy. For the purposes of testing, all redundant components (fabric, routing engines, power supplies, memory cards etc) **should be** present in the system.

Note. For purposes of MPLS VPN forwarding (VPLS, PWE, IPv4/IPv6 VPN), the ports on SUT can be grouped into "network" and "access" side ports according to vendor discretion. Every "network" side port should be configured to send to every "access" side port and vice versa. Every "access" side port should belong to all VPN instances.

All traffic streams on one side (network or access) are required to be of the same capacity. No traffic is allowed between access ports.

Full MAC/route learning in VPN table should be done during the qualification run.

Example 1. A VPLS carrier Ethernet switch has 100x GE ports on "access" side and 10x 10GE ports on the "network" side. Every access port on SUT is divided into 128 VLANs, each VLAN belonging to one of the 128 VPLS instances configured. Router tester simulates no less than one MAC address per each local and remote VLAN.

Every simulated PE sends one stream towards each of the CPE at the same time. Symmetric traffic mesh is created in the opposite direction.

Example 2. An IPv4 VPN is configured on the SUT. Each CPE on the access side sends no less than one route to the VPN, Each remote PE sends no less than one route to the VPN. Every simulated PE sends one stream towards each of the CPE at the same time. Symmetric traffic mesh is created in the opposite direction.

C3.2 Datacenter/Large Enterprise Switching Platforms

Description. Carrier-grade Ethernet switching platforms
Qualification. L2 (Ethernet, MPLS) forwarding, L3 (IPv4, or IPv6 forwarding)
Test Application: L2 or L3 forwarding at vendor discretion. Payload packet size: 256B frames; forwarding over any types of forwarding entries and encapsulation types.
Interface types: SR optical (10/100/1000/10GE) at vendor discretion
Redundancy. For the purposes of testing, all redundant components (fabric, routing engines, power supplies, memory cards etc) **should be** present in the system.

C3.3 Desktop/Generic Ethernet Platforms

Description. Ethernet switching platforms
Qualification. L2 (Ethernet) forwarding, MPLS forwarding, IPv4, or IPv6 forwarding
Test Application: Ethernet or MPLS forwarding at vendor discretion. Payload packet size: 512B frames; forwarding over any types of forwarding entries and encapsulation types.
Interface types: Copper or SR optical (10/100/1000/10GE) at vendor discretion
Redundancy. For the purposes of testing, redundant components *may be* removed

Class 4 - Experimental

Placeholder for any equipment type not assigned to a particular class.
Results in this category can reported along with precise description of the test methodology, topology, traffic profile and protocol configuration.

Class 5 - Security appliances (DPI, Firewalls, VPN Gateways etc)

Description. Security platforms of variable purposes (IP Sec VPN, HTTPS, DPI, IDS etc)
Qualification. L3 forwarding, security features
Test Application: IP SEC or HTTPS, minimum number of firewall or DPI forwarding rules at vendor discretion; 512B payload packets
Interface types: at vendor discretion
Redundancy. For the purposes of testing, redundant components *may be* removed

Class 6 - Application Gateways (Layer 5-7 accelerators, load balancers, etc)

Description. Application platforms of variable purposes (SLB, accelerators, compressors)
Qualification. Application-specific features
Test Application: User traffic at vendor discretion (need more qualification for setup); 512B payload packets
Interface types: at vendor discretion
Redundancy. For the purposes of testing, redundant components *may be* removed

APPENDIX C. FAQ

1. Standardization

Q1.1. – *Standardization* - There are reports of various organizations involved in defining efficiency criteria for network and telecom. Can you clarify the role of ECR in this process?

A. ECR is not a standard, but an open and peer-reviewed packet platform testing methodology and energy efficiency metric designed to be a turnkey solution for national and international standard bodies and organizations.

The following chart demonstrates the worldwide state of network efficiency standartization as of 11/6/2008.

	High-level Definition*	Metric Definition		SUT test procedures	
		Peak	Weighted	Methodology	Load/Profiles
ATIS draft NIPP-TEE-2008-031R4	√	TEER √	not defined	not defined	not defined
METI “Top Runner”	√	not defined	X	P	P
ITU-T♦	not defined	not defined	not defined	not defined	not defined
EC	not defined	not defined	not defined	not defined	not defined
BWF♣	not defined	not defined	not defined	not defined	not defined
Green Grid	√	ATIS	ATIS	ATIS	ATIS
ECR	yes	yes	yes	yes	yes

Legend:

* = Definition of general ICT energy efficiency as “payload per energy unit”
♦ = ECR draft submitted as formal contribution to Climate Change ICT FG
♣ = ECR draft submitted to Marketing and Test & Interoperability/Green OC
√ = identical to ECR
X = currently incompatible with ECR
P = compatible with ECR

Q1.2 – *Standardization* - What is homologation and what agencies are involved?

A. Homologation is conforming equipment to national or international standards. We expect ECR methodology to influence homologation practices in EU (IEC), USA (EPA) and Japan (METI)

2. General

Q2.1 – *General* - why ECR normalizes energy consumption to payload as opposed to using per-port energy allowances defined in [METI 2008] and [EC CoC Broadband] documents?

A. METI "Small Routers" and EC "CoC Broadband" documents were primarily targeted at consumer-level network equipment with limited performance and capacity.

As such, these documents define fixed sets of energy allowances for every product type and functionality bucket. The fact that this option may not operate at line rate typically does not matter in small office/home environment. As a result, consumer-level network and telecom equipment can be massively oversubscribed from the bandwidth perspective without noticeable impact on usability. For instance, it does not matter if the home DSL router cannot operate all wireless or wired LAN ports at line rate, as sustained performance is not required for domestic LAN. Consumer-grade network device can be easily compared to a lightbulb – it fills a basic need at a fixed energy cost.

Carrier-class network and telecom equipment, on the other hand, presents a different case, where functions are delivered across many ports at high speed and revenue generation depends on performance. In the carrier world, an oversubscribed platform is not equal to line-rate device application-wise, and thus, it cannot be fairly compared from the energy consumption perspective.

Currently, the notion of payload-normalized efficiency is universally supported by METI (Top Runner program), ATIS and Green Grid as the basis for telecom/carrier grade efficiency metrics.

Q2.2. - *General* - why would vendors be interested in reporting ECR?

A. Two reasons – first, position their equipment for energy-efficient networking; second – to comply with upcoming government and business standards, such as those expressed by METI in Japan, or ATIS in US.

Q2.3 - *General* - quite obviously, one cannot compare, say, firewall to an optical shelf in energy efficiency. How ECR deals with that?

A. ECR has the notion of product classes – well-defined groups of equipment with similar functionality (core routers, firewalls, broadband aggregation devices etc). The ECR rating has relative significance within the class. For example, the absolute ECR value of a metric for the core-class router is expected to be much better than an enterprise-class router, but they are not comparable to each other.

Q2.4. - *General* - can ECR test methodology be manipulated to achieve better readings?

A. Every test methodology can and should be improved over time – and we see a current example of this in how car manufacturers are forced into using more precise mile-per-gallon test methodology. In the design of ECR, we put significant effort to make sure the system-under-test performance cannot be skewed during the test phase. Since ECR draft is open and peer-reviewed, this also ensures that potential errors or loopholes should be found and fixed with corrections to the test procedure.

Q2.5. - *General* - ECR is a telecom/network device-oriented metric. By focusing on device efficiency, are you ignoring the larger picture**?**

A. Having the best-in-class network components with respect to energy efficiency is not enough to build sustainable or operationally efficient network, but it's a necessary first step. It should be followed with responsible network design (which includes non-telecom infrastructure) and network operation. ECR can be considered a metric for building the bricks. It's possible to build a bad house with good bricks, but if the original goal is to build a good house, it's much easier to do with proper elements.

Q2.6 - *General* - in ECR test procedure, the SUT is equipped up to the maximum. However, in many applications, it won't be the case. Would the measured ECR metric still be relevant?

A. Modular telecom platforms are rarely deployed in full configurations from the start; instead, they typically reach their service ceiling midlife, when the network goes through expansion and upgrade rounds.

To estimate the effect of the 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 linecards, 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.

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

As an illustration, here is an example efficiency profile as collected over an ASIC-based Ethernet switch with 11x linecard slots at all fill levels (one to eleven linecards installed).

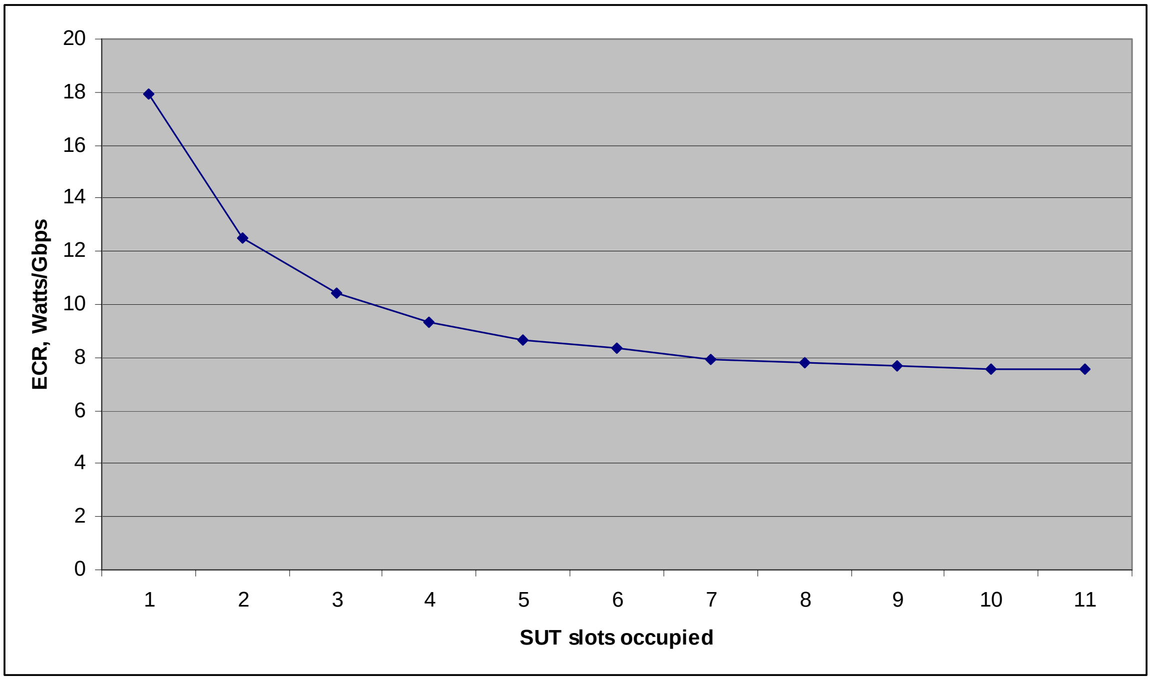

3. Test procedure

Q3.1. – *Test procedure* - how is the SUT probed for effective throughput Tf?

A. We do not define the exact probing and search algorithm for zero-loss operation. We suggest using RFC2544 methodology and applications for doing so – i.e. binary search for correct load profiles. Any other algorithm, including heuristics and well-known performance estimates are accepted.

Q3.2. - *Test procedure* - why is zero-loss operation required? RFC2544 allows for configurable percentage of packet loss.

A. Indeed, there are cases, where application class prevents lossless operation at exact line rate (theoretical physical line limit). Examples would be – interfaces with byte stuffing (i.e. SONET), exception traffic leaving the router/switch (IGP/EGP updates), etc. In this case, the RFC2544 procedure needs to be instructed to top at a safe load level – for instance, 98 or 99 percent theoretical line load. This should not affect relative platform standings as all equipment belonging to the same class would have to be tested in a similar way. On the other hand, random (even minimal) packet loss is very undesirable to modern packet platforms and should be avoided at all costs.

Q3.3 - *Test procedure* - why is the test run defined at 1,200 seconds?

A. We need a compromise between the accuracy and speed. Currently suggested value of 1,200 seconds in most cases will allow full ECR test suite to complete in approximately 1.5 hours. Longer test runs would increase the run time accordingly and risk tying up expensive resources for extended period of time. Shorter test runs may bring the danger of overestimating the SUT. For example, some excess traffic (or traffic bursts) can be wrongfully accounted as delivered, while it could be actually buffered inside the device under test. Also, SUT's energy ratings may be reasonably affected by the state of its active cooling system, which might require a certain temperature threshold to activate (i.e. spin fans at full speed).

Q3.4 - *Test procedure* – why can't SUT be reconfigured between the test runs?

A. This requirement is there to reflect the dynamic nature of internet traffic and associated load profiles. While, indeed, it is often possible to statically alter configuration to match the relaxed load (i.e. remove unused ports, fabric cards, lookup engines, etc), this is not a viable case in the field situation, where the effective load can change at any second. Our EER metric design promotes automatic (intelligent) power management schemes.

Q3.5 - *Test procedure* - what determines the order of test runs and the time between runs?

A. The order of runs can affect energy readings because electronic platforms may change their energy consumption based on the component temperature which, in turn, can be driven by utilization. This means that energy readings collected for 100 percent, 50 percent and idle runs will not be the same depending on the order and interval. On the other hand, some idle intervals are needed for most packet testers (ramp-on/ramp-off statistics generation), which can be followed with SUT "priming" for state- or activity-driven protocols.

Q3.6 - *Test procedure* – why is packet size variable between product classes?

A. With respect to the packet size, the SUT may have two responses – (a) performance limits, (b) energy consumption change due to changes in packet lookup budgets/cycles. An example of energy-to-packet size graph is given below:

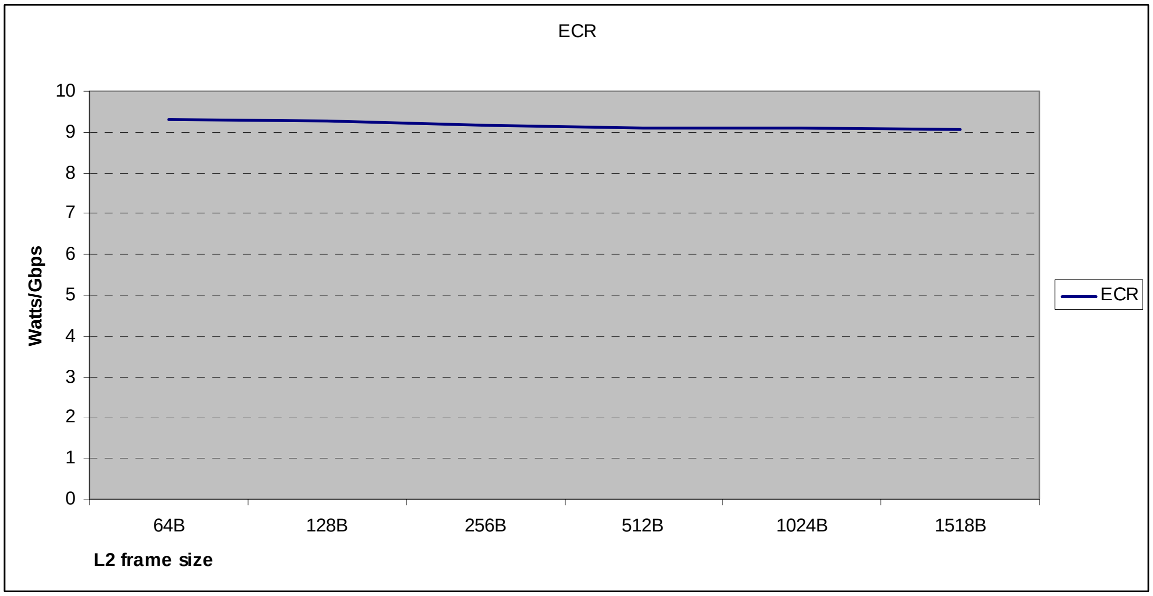

While it could be safer to measure every SUT against an array of packet sizes, this would result in explosion of test runs (N packet sizes x N load levels). Therefore, we decided to assign each class with a fixed packet size commensurate with application.

Q3.6 - *Test procedure* – why ECR does not use packet mixes?

A. Certain packet mixes (i.e. different iMix compositions) are known to exercise the SUT datapath in the more realistic format than uniform packet streams. However, the lack of the agreed-upon iMix generation and reporting (statistical vs. fixed sequence) across packet / router tester vendors at this moment is holding us back from using packet mixes.

In addition, the implication of statistical load compared to measured system performance at this moment is not completely understood. This section may be revisited in the future.

4. Metric

Q4.1 – *Metric* - what is the difference between synthetic and physical-meaning metrics?

A. Metrics with physical meaning can be attributed to the actual physical process (under controlled conditions). For example, cars are rated in miles-per-gallon, which reflects their ability to transport passengers or cargo under the rules established by US EPA. Likewise, ECR is a peak metric, expressed in Watts/Gbps. This reflects the amount of energy (in Joules) required to move a fixed amount of data (in Gigabits) under specific conditions (configuration and load profile).

On the other hand, any metric that cannot be related to a physical process, is synthetic and cannot be expressed in physical units. Consider the following example:

A SUT is being evaluated against a metric, which uses weights to rate the relative importance of load profiles: $Tf/(a*Ef + b*Eh + c*Ei)$, where Tf = measured throughput at full load, { Ef, Eh, Ei } – energy consumption at full, half and idle load respectively.

If SUT draws 1,000 Watts at utilization 100Gbps, 100 Watts at utilization 10Gbps and 50 Watts at idle, and weights {a, b, c} are given the values of a=0.1, b=0.6, c=0.3,

The metric becomes: $100G/(0.1 * 1000W + 0.6 * 100W + 0.3 * 50W) = 0.57$

Although the source data was expressed in Gbps and Watts, the resulting metric cannot be given a “Gbps/Watt” designation (consider the peak efficiency 100/1,000 = 0.1 Gbps/W)

Q4.2 – *Metric* - Can ECR/ECRW be used to compare hybrid devices?

A. Absolutely. Hybrid devices are the systems that patently can operate in different equipment classes – i.e. a router and a firewall. We recommend vendors of hybrid devices to obtain the ECR metrics in all relevant classes the device can be certified, for example ECR(Class A) and ECR(Class B). In this case, a hybrid device can be compared to a stack of single-purpose devices according to the user-defined proportion of functionality required in the respective classes (ECR ratings are additive).

Q4.3 – *Metric* – Is there a specific measurement mode for “green” or “sleep” states?

A. ECR methodology already incorporates the dynamic power management ability of SUT in ECRW metric, which is a supplementary to peak ECR efficiency. If SUT is capable of reducing energy consumption at reduced load, this would result in better ECRW, which is very relevant to the dynamic nature of real-life networks

If SUT is also capable of reducing energy consumption based on capacity planning and configuration changes (static energy management capability), this should designate a new “platform”, where a separate SUT entry is created with it’s own ECR and ECRW metrics.

For example, “SUT A” and SUT “A-Green mode” are likely to have different performance and energy consumption depending on the class of usage/load profile.

5. Reporting

Q5.1. *Reporting* – is there a specific format the full ECR results can be reported with?

A. The exact format does not exist at this time, but we expect it to emerge over time from standard bodies, such as ATIS.