

Network and Telecom Equipment— Energy and Performance Assessment

Metrics, Test Procedure and Measurement Methodology

Draft 3.0.1, December 14, 2010



Revision Record:

Version	Date	Author	Reason
0.2	12/12/2007	dkh	First draft
0.3	03/22/2008	dkh	First version of weighted metric
0.5	06/11/2008	aalimian	Extended product classes
0.6	06/15/2008	dkh	Changes per first ECR trials across different product types
0.9	07/15/2008	dkh	Changes and comments per input from Bruce Nordman (LBNL)
1.0	10/12/2008	dkh	Changes per VPLS test trials
1.0.3	11/06/2008	dkh	Update to product classes
1.0.6	02/18/2009	dkh	Clarifications to basic test topologies T1 and T2
1.1.0	04/10/2009	dkh	Section on cascaded energy management
1.2.1	06/01/2009	dkh	Separate test procedure for idle (static) port utilization mode
1.4.0	08/01/2009	dkh	Separate test procedure for energy usage self-monitoring
1.5.0	09/22/2009	dkh	Section on cascaded energy management
2.0.1	09/22/2009	dkh	Migration to a second version of weighted metric (ECR-VL)
3.0.1	12/04/2010	dkh	Extended idle metric addition (ECR-EX)

Areg Alimian, IXIA Corp
Bruce Nordman, Lawrence Berkeley National Lab
Daniel Kharitonov, Juniper Networks, Inc.

aalimian@ixiacom.com
BNordman@lbl.gov
dkh@juniper.net

Table of Contents

Copyright Notice	4
List of Acronyms	4
Purpose	5
Theoretical Basis	5
Scope	5
Consistency	6
Result Representation and Applicability	6
Base Metrics and Device Comparisons	6
Energy Bill Estimates	9
Site Planning	11
Functional Compliance	12
Reporting Format	13
Product Labeling	13
APPENDIX A: Measurement Methodology	14
Test Procedure 1 (Mandatory)	14
Test Procedure 2 (Optional)	16
Test Procedure 3 (Optional)	17
Test Procedure 4 (Optional Functional Test)	19
Test Procedure 5 (Optional Functional Test)	22
Effective Throughput Calculation	22
APPENDIX B: Measurement Conditions	24
APPENDIX C: Proposed Product Classes and Test Applications	25
Topologies in Use	25
Class 1—Routers	26
C1.1 Core Routers	26
C1.2 Carrier Edge Routers/Ethernet Service Routers	27
C1.3 Multipurpose Routers	27
Class 2—WAN/Broadband Aggregation Device	28
C2.1 BRAS Devices	28
C2.2 BNG/Common Edge devices	28
Class 3—Ethernet L2/L3 Switches	29
C3.1 Core/Datacenter Ethernet Switching Platforms	29
C3.2 Desktop/Aggregation Ethernet Platforms	29
Class 4—Experimental	29
Class 5—Security Appliances (DPI, Firewalls, VPN Gateways)	29
Class 6—Application Gateways (Layers 5 Through 7 Accelerators, Load Balancers)	30
References	31

Copyright Notice



This text, “Network and Telecom Equipment—Energy and Performance Assessment,” is licensed under a Creative Commons Attribution 3.0 United States License. This license lets others distribute, remix, tweak, and build upon this work, even commercially, as long as they credit this original document.

List of Acronyms

ECR	energy consumption rating (W/Gbps)
ECR-VL	energy efficiency metric over a variable-load cycle (W/Gbps)
ECR-EX	energy efficiency metric over extended-idle load cycle (W/Gbps)
L_{max}	maximum offered load at zero packet loss
T_f	measured maximum effective throughput (full-duplex, Gbps)
E₁₀₀	energy consumption under highest load accepted by SUT (watts)
E₅₀	energy consumption under fifty percent offered load (watts)
E₃₀	energy consumption under thirty percent offered load (watts)
E₁₀	energy consumption under ten percent offered load (watts)
P₅₀	energy consumption with half capacity active (watts)
P₂₅	energy consumption with one-quarter capacity active (watts)
P₁₀	energy consumption with one-tenth capacity active (watts)
R₁₀₀	energy consumption under L_{max} as reported by SUT (watts)
U₁₀₀	reported system utilization under load L_{max} (percent)
SUT	system under test
CUT	component under test
F_c	Energy footprint of a component (estimated, watts)

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. Various aspects of operation are covered, including peak efficiency, variable-load efficiency and idle (statically configurable) energy efficiency.

Portions of this document can be used free of charge under the terms of Creative Commons Attribution 3.0 United States License. (See copyright notice on previous page.)

Theoretical Basis

This document defines energy efficiency as energy consumption normalized to effective throughput. Such approach is in accordance with high-level methodology suggested in [SAINT 2008] and similar documents. In other words, we assume the more energy-efficient network system to be the one that can transport more data (in bits) using the same energy budget (in Joules).

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

Because the amount of data the system can transport may be different from the highest theoretically possible load, system performance should be verified alongside with energy consumption.

In addition to peak energy performance, separate test profiles are used to measure energy response to variable-load (realtime) and extended idle (non-realtime) utilization profiles.

Additional (optional) tests can be added for comprehensive evaluation of other energy-related properties, such as energy management for connected (cascaded) devices and embedded energy monitoring capabilities.

Scope

Efficiency definition provided above is best suited for medium- to large-scale network and telecom systems. It is less relevant to small office, home office and consumer-grade multipurpose communication devices, where throughput is less relevant and energy efficiency metrics can be based on allowances per units of functionality, such as described in [METI 2008] and [EC CoC Broadband] documents.

In general, this document is applicable to many types of packet-oriented network and telecom equipment, including, but not limited to, core and edge routers, L2/L3 switches, packet-optical equipment, security devices, load balancers, etc.—anything that can exhibit performance numbers lower than the face value of connected ports.

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

Consistency

This test methodology is designed to produce results consistent with actual SUT capabilities and free of misrepresentation. It is recommended to accept ECR test results only when certified by the 3rd party.

Result Representation and Applicability

The array of results obtained in Test Procedures 1 through 4 forms the energy “passport” of the SUT and can be used directly by consumers for evaluation and energy planning purposes.

Base Metrics and Device Comparisons

Comparing product metrics allows consumers, enterprises and carriers to add energy efficiency to purchase criteria. The most straightforward way to estimate the technology level of a network or telecom system is to normalize its energy consumption to the highest sustained throughput recorded in the test.

$$ECR = E_{100}/T_f \quad (\text{expressed in } W/Gbps)$$

where

T_f = maximum throughput (Gbps) (See “[Effective Throughput Calculation](#)”)

E_{100} = energy consumption (watts) measured during step 2 of Test Procedure 1

ECR¹ (energy consumption rating) is normalized to W/Gbps and has a physical meaning of energy consumption to move one gigabit worth of line-level data per second. ECR typically reflects the best possible platform performance within a set of hardware and software features².

Although ECR is an accurate measure of SUT technology level, network systems in the field are unlikely to demonstrate comparable efficiency numbers over sustained intervals. This situation is related to the fact that packet networks tend to exhibit low long-term utilization numbers coupled with significant short-term bursts (peaks), which causes service providers to size network infrastructure based on the higher end of traffic profiles and subsequently lose energy efficiency during off-peak times.

There is compelling opportunity for energy savings here because vendors can choose to optimize their products for energy consumption in the middle—rather than at the top—of

¹ Energy efficiency can also be reported in Gbps/W (EER); $EER = 1 / (ECR)$

² Other ECR denominations can also be used, such as W/10 Gbps, W/100 Gbps or Joules/Gb.

their operational load band. This fact can be reflected in a weighted, variable-load metric, such as ECR-VL:

$$\text{ECR-VL} = \frac{(\alpha * E_{100} + \beta * E_{50} + \gamma * E_{30} + \delta * E_{10} + \epsilon * E_i)}{(\alpha * T_f + \beta * T_{50} + \gamma * T_{30} + \delta * T_{10})}$$

where

T_f = maximum throughput (Gbps) achieved in the measurement cycle

$T_{50} = T_f * 0.5$

$T_{30} = T_f * 0.3$

$T_{10} = T_f * 0.1$

E_{100} = energy consumption (watts) measured during step 2, Test Procedure 1

E_{50} = energy consumption (watts) measured during step 3, Test Procedure 1

E_{30} = energy consumption (watts) measured during step 4, Test Procedure 1

E_{10} = energy consumption (watts) measured during step 5, Test Procedure 1

E_i = energy consumption (watts) measured during step 6, Test Procedure 1

$\alpha, \beta, \gamma, \delta, \epsilon$ are weight coefficients selected such as $(\alpha + \beta + \gamma + \delta + \epsilon) = 1$

ECR-VL is measured in W/Gbps and has a physical meaning of an average energy rating in a reference network described by array of utilization weights ($\alpha, \beta, \gamma, \delta, \epsilon$).

In this weight array, α represents a relative weight of a peak-utilization period and ϵ represents a relative weight of an idle interval during which a network system does no useful work. All other weights are sized for relative duration of the intervals in between.

An ideal network system following the Barroso's principle of energy-proportional computing [IEEE Computer 2007] should be able to demonstrate an ECR-VL rating to be close to ECR. Therefore, ECR-VL is a measure of dynamic (real-time) energy management capability of a network device³.

Network environments with well-known time cycles (such as enterprises) may achieve additional energy savings by enabling their network devices to enter non-realtime power states during periods of low utilization. Such capability represents the tradeoff between energy savings and possible (temporary) traffic loss during unexpected traffic surges.

This fact can be reflected in a weighted, extended-idle metric, such as ECR-EX:

³ Although no real network can be adequately described with a scalar combination of weights, a simplified network-utilization profile can still result in a metric suitable to optimize network and telecom products for load-proportional energy consumption.

$$\text{ECR-EX} = \frac{(\alpha * E_{100} + \beta * P_{50} + \gamma * P_{30} + \delta * P_{10})}{(\alpha * T_f + \beta * T_{50} + \gamma * T_{30} + \delta * T_{10})}$$

where

T_f = maximum throughput (Gbps) achieved in the measurement cycle

$T_{50} = T_f * 0.5$

$T_{30} = T_f * 0.3$

$T_{10} = T_f * 0.1$

E_{100} = energy consumption (watts) measured during step 2, Test Procedure 1

P_{50} = energy consumption (watts) measured during step 1, Test Procedure 2

P_{30} = energy consumption (watts) measured during step 2, Test Procedure 2

P_{10} = energy consumption (watts) measured during step 3, Test Procedure 2

$\alpha, \beta, \gamma, \delta, \varepsilon$ are weight coefficients selected such as $(\alpha + \beta + \gamma + \delta + \varepsilon) = 1$

ECR-EX is measured in W/Gbps and has a physical meaning of an average energy rating in a reference network, where non-realtime (extended) energy savings capabilities are enabled.

Example 1

The SUT has demonstrated the following results:

$T_f = 627.20 \text{ Gbps @ } E_{100} = 5,856 \text{ Watts}$

$T_{50} = 313.60 \text{ Gbps @ } E_{50} = 5,616 \text{ Watts @ } P_{50} = 5,120 \text{ Watts}$

$T_{30} = 188.16 \text{ Gbps @ } E_{30} = 5,520 \text{ Watts @ } P_{30} = 4,810 \text{ Watts}$

$T_i = 0 \text{ Gbps @ } E_i = 5,376 \text{ Watts}$

$ECR = E_{100}/T_f = 5,856/627.20 = 9.34 \text{ W/Gbps}$

Considering $\alpha = 0.1, \beta = 0.5, \gamma = 0.3, \delta = 0, \varepsilon = 0.1$:

$ECR-VL = (\alpha * E_{100} + \beta * E_{50} + \gamma * E_{30} + \varepsilon * E_i) / (\alpha * T_f + \beta * T_{50} + \gamma * T_{30}) =$
 $= (0.1 * 5,856 + 0.5 * 5,616 + 0.3 * 5,520 + 0.1 * 5,376) / (0.1 * 627.2 + 0.5 * 313.6 +$
 $0.3 * 188.2) = 5587.2 / 275.968 = 20.25 \text{ W/Gbps}$

$ECR-EX = (\alpha * E_{100} + \beta * P_{50} + \gamma * P_{30}) / (\alpha * T_f + \beta * T_{50} + \gamma * T_{30}) =$
 $= (0.1 * 5,856 + 0.5 * 5,120 + 0.3 * 4,810) / (0.1 * 627.2 + 0.5 * 313.6 + 0.3 * 188.2) =$
 $5587.2 / 275.968 = 16.62 \text{ W/Gbps}$

Example 2

Four routers from different vendors demonstrated the following results:

	Product A	Product B	Product C	Product D
Product class	Core router	Core router	Core router	Core router
Nominal capacity (vendor rated, half-duplex)	640 Gbps	1.28 Tbps	1.6 Tbps	3.2 Tbps
T_f (measured)	300.1 Gbps	627.2 Gbps	790 Gbps	1.42 Tbps
E_{100} (measured)	4,501 W	5,856 W	7,990 W	11,360 W
ECR	15 W/Gbps	9.34 W/Gbps	10 W/Gbps	8 W/Gbps
ECR-VL	30 W/Gbps	20 W/Gbps	12 W/Gbps	14 W/Gbps

From this table we can see the effect of normalization: Although Product A has the lowest energy consumption per system, its relative efficiency numbers are the worst in its class. Likewise, Product D has the highest absolute consumption but demonstrates the best peak efficiency. However, the lowest cost of operation will probably (depending on traffic profile) be shown by Product C, which has the best dynamic-energy management capabilities.

Energy Bill Estimates

It is also interesting to have an estimate of energy consumption over a projected lifetime (cost of operation). For a reference system described by a vector of variable-load measurements (E_{100} , E_{50} , E_{30} , E_{10} , E_i), building such an estimate is trivial and can be expressed by the following equation:

$$C = \sum_{j=1}^N (\alpha * E_{100} + \beta * E_{50} + \gamma * E_{30} + \delta * E_{10} + \varepsilon * E_i) / 1000 * 8765.25 * C_{kwh j}$$

Weights (α , β , γ , δ , ε) are specific to a customer's network, N is a projected lifetime (in years) and $C_{kwh j}$ represents the cost of kilowatt hours in a year of operation.

Example 1

Verizon Network Equipment Building Systems (NEBS) compliance document [VZ.TPR.9205] estimates its network utilization weights to be ($\alpha = 0.35$, $\beta = 0.4$, $\gamma=0$, $\delta=0$, $\varepsilon = 0.25$). Assuming a cost of a kilowatt hour to be \$0.1 in the first year of operation (and rising one cent every year), the following piece of equipment (LAN switch) would have the following cost of operation over five years:

$$T_f = 627.20 \text{ Gbps @ } E_{100} = 5,856 \text{ W}$$

$$T_{50} = 313.60 \text{ Gbps @ } E_{50} = 5,616 \text{ W}$$

$$T_{30} = 188.16 \text{ Gbps @ } E_{30} = 5,520 \text{ W}$$

$$T_i = 0 \text{ Gbps @ } E_i = 5,376 \text{ W}$$

$$C = \sum_{j=1}^5 (\alpha * E_{100} + \beta * E_{50} + \epsilon * E_i) / 1000 * 8765.25 * C_{kwh j}$$

$$C = (0.35 * 5,856 + 0.4 * 5,616 + 0.25 * 5,376) / 1000 * 8760.25 * (0.1 + 0.11 + 0.12 + 0.13 + 0.14)$$

$$C = (0.35 * 5,856 + 0.4 * 5,616 + 0.25 * 5,376) / 1000 * 8760.25 * (0.1 + 0.11 + 0.12 + 0.13 + 0.14)$$

$$C = (2,050 + 2,246 + 1,344) * 8760.25 * 0.6 = \mathbf{\$29,645}$$

This method gives the estimate based on automatic (variable-load) energy management capabilities of a reference system.

A different kind of estimate can be built upon static (idle-load) energy savings capabilities of a network device (if present). For instance, if a LAN switch in the enterprise can be safely degraded to 25% of its port capacity during the night time (00:00h to 6:00h) and down to 10% during weekends and national holidays, the formula for lifetime energy cost can be expressed as:

$$C = \sum_{j=1}^N (a * P_{100} + b * P_{25} + c * P_{10}) / 1000 * C_{kwh j}$$

where

- a = number of workdays per year * number of work hours per day
- b = number of workdays per year * idle (off-shift) hours per workday
- c = number of holidays per year * 24 hours

For most US locations, the following numbers are applicable:
a = 261 * 18 = 4,698 hours; b = 261 * 6 = 1,566 hours; c = 104 * 24 = 2,496 hours

Example 2

A LAN switch is located in the U.S. and consumes 5,856 W at full load with all ports passing traffic. It can also be configured for energy-saving operation with 25% port capacity ($P_{25} = 5,201$ W) and 10% port capacity ($P_{10} = 4,225$ W). Assuming the 25% port utilization policy is in place from 00:00h to 06:00h every workday and a 10% policy is in place every weekend and holiday, energy costs over five years will be:

$$C = \sum_{j=1}^N (a * P_{100} + b * P_{25} + c * P_{10}) / 1000 * C_{kwh j}$$

$$C = \sum_{j=1}^5 (4,698 * 5,856 + 1,566 * 5,201 + 2,496 * 4,225) / 1000 * C_{kwh j}$$

Assuming the cost of a kilowatt hour to be \$0.1 in the first year of operation and rising one cent every year:

$$C = (4,698 * 5,856 + 1,566 * 5,201 + 2,496 * 4,225) / 1000 * (0.1 + 0.11 + 0.12 + 0.13 + 0.14)$$

$$C = (27,511,488 + 8,144,766 + 10,545,600) / 1000 * 0.6 = \$27,721$$

Although it might look like the cost of operation with active idle-load management (Example 2) is lower compared to the variable-load capabilities built into the SUT (Example 1), in practice, both methods are complementary and operate on different time intervals.

Variable-load energy management capabilities are typically modest, but they provide lossless system operation under all conditions.

Idle-load management is based on non-realtime transition between power states and might not be applicable to environments with unpredictable utilization peaks. Static energy management also requires network support personnel to take responsibility for policy-driven capacity control and possible service degradation if the network experiences unexpected utilization spikes.

For more discussion on relation of silicon technologies and time intervals, please refer to publication [Globecomm 2009].

Site Planning

When planning main and backup power, cooling, and operational budgets for telecom points of presence, network engineers traditionally use power ratings (agency labels) provided by network equipment vendors. In practice, power ratings result in conservative site planning because they call for energy reservations at the highest end of a possible spectrum. This creates a common problem when onsite power capabilities are below the agency-label levels and it is unclear whether they can support a newly proposed network system in a specific (even lightly loaded) configuration (for more discussion on rated versus measured energy ratings, see publication [JNPR TSBG]).

The following resources are useful for purposes of site planning:

- Value E_{100} , which describes the average SUT energy consumption under highest possible load. When coupled with safety margins, this value can be used as an upper boundary for energy requirements of a reference system with components of a same or similar type.
- When a vendor does not provide E_{100} measurements for a specific SUT configuration of interest, this value can be approximated from a sum of required components (provided that component footprints F_{ci} were published).

Example 1

A deep-packet inspection system is rated at 4550 W of maximum energy consumption rate. The same device was measured at $E_{100} = 3400$ W in a reference configuration.

Question: Can this equipment be installed in a cabinet with 4,000 W of power and cooling capacity?

Answer: Yes, unless a vendor says otherwise. The risk here is that the future generations of hardware available for the same system might not fit into the currently measured energy budget E_{100} .

Example 2

A modular edge router is measured at $E_{100} = 4,900$ W in a reference configuration with all slots filled by line cards (10 slots x quad-port 10GbE cards).

Question: What is the energy budget for the same system with only four slots occupied by the same quad-port 10GbE linecards?

Answer: This question can be answered with “Test Procedure 3”. For instance, if a full system consists of five component types, such as base (C1), fabric (C2), routing engine (C3) and line cards (C4) and vendor provided ratings F_{c1} F_{c2} F_{c3} F_{c4} , the approximate system energy footprint F_{100} might look like this:

$$F_{100} = (F_{c1} + F_{c2} + F_{c3} + 4 * F_{c4})$$

Functional Compliance

Test Procedures 4 and 5 are intended for functional compliance testing with pass/fail criteria. This functionality can be used to understand, plan for and manage the energy consumption within a telecom network and around it, both at the enterprise (LAN, data center) and consumer (home LAN, “Smart Home”) levels.

In Test Procedure 4 (embedded energy-monitoring capabilities), the externally measured energy consumption values (E_{100} , P_{50} , P_{25} , P_{10}) should be compared to values (R_{100} , R_{50} , R_{25} , R_{10}), reported by the SUT. The adjusted result should correlate to effective consumption with acceptable precision ($\pm 10\%$ or less) across all measurement points; Upon vendor approval, a static offset can be added to R to compensate for electronics not covered by embedded power monitoring circuitry.

Comprehensive, device-level energy consumption monitoring across the entire telecom network help to identify opportunities for operational energy savings. Likewise, externally recorded load L in the test (where $L_{max} = 100\%$ SUT utilization) should correlate to values U reported by the SUT. Upon vendor approval, a static offset can be added or subtracted to adjust for load calculation algorithms employed by the SUT.

The knowledge of a system with respect to the timeline is critical for planning the static energy management routines (see “Energy Bill Estimates”). Test Procedure 5 describes the value-add functionally that might be present in the SUT.

APPENDIX A: Measurement Methodology

Test Procedure 1 (Mandatory)

Energy consumption in relation to variable load

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. Additional class-specific requirements for this test are listed in Appendix B.

Note: There is no SUT configuration change allowed any time beyond the preparation phase. All energy savings adjustments (if done) by the SUT during Test Procedure 1 should be automatic.

SUT Preparation

The SUT is configured according to class requirements and offered the load as defined in the class requirements (Appendix B). Typically, the SUT is a reference system fully outfitted with hardware and software suitable for class requirements. 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 and humidity differences.

Router testing equipment (traffic generators) 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. Metrology and environmental requirements to DC- and AC-based test cases are listed in Appendix A.

This measurement procedure consists of six major steps.

Step 1 (qualification)

The first run determines the maximum load (L_{\max}) 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.

Note: The following test runs should be separated with an 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 in the idle-time window.

Step 2 (full load)

The second run offers the load L_{\max} (identified at step 1) to the SUT for a period of 1200 seconds⁴. Energy consumption is being sampled for the entire period and average consumption E_{100} calculated⁵.

Step 3 (half load)

The third run reduces the load L_{\max} twice ($L_{\text{half}} = 0.5 \times L_{\max}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption E_{50} calculated. Load reduction is achieved by *reducing the packet rate* on all configured ports. (Load reduction by means of idling or disconnecting ports is *not* acceptable.)

Packet loss during this run (if seen) invalidates the measurement and resets testing to the first run to provide a better L_{\max} estimate.

Note: If the SUT does any energy management to adapt to the lower offered rate, this action should be fully automatic and should allow instant return to the offered rate L_{\max} without packet loss. Random probing for higher load levels is recommended to validate test results.

Step 4 (30% load)

This test run further reduces the load to L_{30} ($L_{30} = 0.3 \times L_{\max}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption E_{30} calculated. Load reduction is achieved by *reducing the packet rate* on all configured ports. (Load reduction by means of idling or disconnecting ports is *not* acceptable.)

Packet loss during this run (if seen) invalidates the measurement and resets testing to the first run to provide a better L_{\max} estimate.

Note: If the SUT does any energy management to adapt to the lower offered rate, this action should be fully automatic and should allow instant return to the offered rate L_{\max} without packet loss. Random probing for higher load levels is recommended to validate test results.

Step 5 (10% load)

This test run further reduces the load to L_{10} ($L_{10} = 0.1 \times L_{\max}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption E_{10} calculated. Load reduction is achieved by *reducing the packet rate* on all

⁴ The measurement interval represents a compromise between accuracy and speed. Shorter time intervals (i.e., 300 seconds) might 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.

⁵ Please refer to Appendix A for measurement conditions and qualifications.

configured ports. (Load reduction by means of idling or disconnecting ports is *not* acceptable.)

Packet loss during this run (if seen) invalidates the measurement and resets testing to first run to provide a better L_{\max} estimate.

Note: If the SUT does any energy management to adapt to the lower offered rate, this action should be fully automatic and should allow instant return to the offered rate L_{\max} without packet loss. Random probing for higher load levels is recommended to validate test results.

Step 6 (idle run)

Idle run removes the load and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption E_i calculated. Load reduction is achieved by idling packet rate on all configured ports. (Load reduction by means of disconnecting or shutting down ports is *not* acceptable.)

Note: If the SUT does any energy management to adapt to the lower offered rate, this action should be fully automatic and should allow instant return to the offered rate L_{\max} without packet loss. Random probing for higher load levels is recommended to validate test results.

Test Procedure 2 (Optional)

Energy consumption in relation to static load change

In addition to dynamic energy performance demonstrated at peak and variable offered load, network and telecom systems might possess static energy management capabilities, adapting to extended periods of low utilization by entering various power conservation states. This measurement procedure intends to demonstrate the energy savings potential of such features (if present).

The prerequisite for this procedure is to have a value L_{\max} determined during Test Procedure 1. The SUT configuration, preparation and class-specific requirements, and test bed setup for this procedure are identical to Test Procedure 1. This measurement procedure consists of three major steps.

Note: The following test runs should be separated with an idle time of 300 seconds or less. It is acceptable to change the SUT configuration to move components into a low-power or shutdown state prior to each step. This should be done in the idle-time window.

Step 1 (half capacity in use)

This test run reduces the load L_{\max} twice ($L_{\text{half}} = 0.5 \times L_{\max}$) and runs for 1200 seconds. Energy consumption is being measured for the entire period and average consumption P_{50} calculated. Load reduction is achieved by either reducing traffic rate across all ports or by *sending traffic at full rate to every second port of the SUT*. Unused router tester ports can

be left idle or shut down at vendor discretion. Active ports should be evenly mixed with inactive ports on the SUT; block designation (i.e., four active ports on one linecard, then four inactive ports on another linecard) is *not* acceptable.

Note: Packet loss during this run (if seen) invalidates the measurement and resets testing to Test Procedure 1 to provide a better L_{\max} estimate. In this test, SUT is *not* required to return to higher performance levels instantly and *might* require a configuration change and recovery period to return to full capacity.

Step 2 (one quarter capacity in use)

This test run further reduces the load to L_{25} ($L_{25} = 0.25 \times L_{\max}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption P_{25} calculated. Load reduction is achieved by either reducing traffic rate across all ports or by *sending traffic at full rate to every fourth port of the SUT*. Unused router tester ports can be left idle or shut down at vendor discretion. Active ports should be evenly mixed with inactive ports on SUT; block designation (i.e., 4 active ports on one linecard, then 12 inactive ports on other linecards) is *not* acceptable.

Note: Packet loss during this run (if seen) invalidates the measurement and resets testing to Test Procedure 1 to provide a better L_{\max} estimate. In this test, SUT is *not* required to return to higher performance levels instantly and *might* require a configuration change and recovery period to return to full capacity.

Step 3 (one-tenth capacity in use)

This test run further reduces the load to L_{10} ($L_{10} = 0.1 \times L_{\max}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption P_{10} calculated. Load reduction is achieved by either reducing traffic rate across all ports or by *sending traffic at full rate to every tenth port of the SUT*. Unused router tester ports can be left idle or shut down at vendor discretion. Active ports should be evenly mixed with inactive ports on SUT; block designation (i.e., 10 active ports across two linecards, and then 90 inactive ports across remaining linecards) is *not* acceptable.

Note: Packet loss during this run (if seen) invalidates the measurement and resets testing to Test Procedure 1 to provide a better L_{\max} estimate. In this test, SUT is *not* required to return to higher performance levels instantly and *might* require a configuration change and recovery period to return to full capacity.

Note: There is no zero-load test because it is assumed that unused equipment can be turned on and off with internal timer/scheduler logic or remotely.

Test Procedure 3 (Optional)

Component-level energy footprint (F_c)

Modular network and telecom systems can be configured in a variety of ways, which renders full testing of all possible reference configurations impractical. Because of this,

vendors can choose to list only one or several (most typical) configurations of the same modular system for detailed energy and performance reports.

Other configurations can be approximated based on a sum of component-level energy budgets. Intermodule dependencies and interactions make such approximation significantly less precise compared to the actual measurement done on a live device, yet this trade-off can be acceptable for operational cost or facility planning purposes.

The following measurement procedure allows a vendor to build and publish a library of component-level energy consumption suitable for power budget estimates of arbitrary complex configurations.

Note: The simplistic nature of such representation means the detailed measurements (i.e., subpeak load) are not applicable. This procedure concentrates on obtaining peak energy performance results only.

SUT Preparation

A modular SUT is configured with hardware, software and connection topology relevant for its class of operation (see Appendix B), including the component (or module) under test (CUT). The SUT is *not* required to be fully loaded and is *not* required to bear all modules of the same type. However, it is beneficial to make sure that the energy and traffic impact of the CUT is minimized compared to the rest of the system, which is best achieved in higher-end configurations.

Step 1 (qualification)

The first run determines the maximum load ($L_{\max+}$) that the SUT can sustain at zero packet loss. It is recommended to design $L_{\max+}$ to exercise all SUT components (including CUT) to the highest possible throughput. A failure to do so might significantly affect test results.

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.

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

Step 2 (full load)

The second run offers the load $L_{\max+}$ (identified in step 1) to the SUT for a period of 1200 seconds⁶. Energy consumption is being sampled for the entire period and average system footprint F_+ calculated⁷.

⁶ The measurement interval represents a compromise between accuracy and speed. Shorter time intervals (i.e., 300 seconds) might 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.

⁷ Please refer to Appendix A for measurement conditions and qualifications

Note: Packet loss during this run (if seen) invalidates the measurement and resets testing to first run to provide a better $L_{\max +}$ estimate.

Step 3 (CUT removal)

The CUT should be removed from the SUT at this step. The SUT and test system configurations might need to be readjusted to compensate for component removal.

The third run determines the new maximum load ($L_{\max -}$) the SUT can be sustained at zero packet loss. It is recommended to design $L_{\max -}$ to exercise all remaining SUT components to the highest possible throughput. A failure to do so can significantly affect test results.

Step 4 (final measurement)

At this step, the load $L_{\max -}$ (identified in step 3) is offered to the SUT for a period of 1200 seconds. Energy consumption is being sampled for the entire period and average system footprint F_c calculated.

Note: Packet loss during this run (if seen) invalidates the measurement and resets testing to the first run to provide a better $L_{\max -}$ estimate.

At this point, the energy footprint of the CUT can be approximated:

$$F_c = F_+ - F_-$$

The same procedure (Steps 1 to 4) can be repeated as many times as needed to build a component footprint library.

Note: This method is not applicable for nonredundant components in the critical system path. Because of this limitation, it is acceptable for some components (e.g., system chassis and backplane) to be reported as “base system” aggregate when their individual footprints cannot be reliably determined.

Test Procedure 4 (Optional Functional Test)

Embedded energy monitoring capabilities

Real-time energy consumption estimates are pivotal to the designing and monitoring of energy conservation policies in telecom points of presence and data centers. In the meanwhile, telecom and datacenter power facilities typically do not provide energy utilization reports with device-level granularity. Therefore, it is recommended for a telecom or network system to have the ability to monitor and record its own energy consumption with embedded probes and sensors. It is also useful to report energy consumption in parallel with system utilization, which allows for traffic affinity and energy pattern analysis.

Note: For time and cost-reduction purposes, this test procedure can be exercised in parallel with “Test Procedure 1” and “Test Procedure 2”.

If there is a need to enable or reset embedded energy management and system utilization agents or counters on the SUT, this should be done prior to step 1.

Step 1 (qualification) – this step can be combined with “Test procedure 1, step 1”

The first run determines the maximum load (L_{max}) 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.

Note: The following test runs should be separated with an 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 in the idle-time window.

Step 2 (full load) – this step can be also combined with “Test procedure 1, step 2”

The second run offers the load L_{max} (identified at step 1) to the SUT for a period of 1200 seconds⁸. Energy consumption is being sampled for the entire period and average consumption E_{100} calculated⁹.

Average recorded energy consumption R_{100} for the same 1200 second interval should be read from the SUT at the end of this run with appropriate means (SNMP, DMI (Device Management Interface), XML agent, command line interface, etc.).

Recommended: Average recorded system utilization U_{100} for the same 1200 second interval should be read from the SUT at the end of this run with the appropriate means (SNMP, DMI, XML agent, command line interface, etc.).

Step 3 (half load) – this step can be also combined with “Test procedure 2, step 1”

Note: It is acceptable to change the SUT configuration to move components into a low-power or shutdown state prior to each step from “Step 3” and beyond. This should be done during the idle-time window.

This test run reduces the load L_{max} twice ($L_{half} = 0.5 \times L_{max}$) and runs for 1200 seconds. Energy consumption is being measured for the entire period and average consumption P_{50} calculated. Load reduction is achieved by either reducing traffic rate across all ports or by *sending traffic at full rate to every second port of the SUT*. Unused router tester ports can be left idle or shut down at vendor discretion. Active ports should be evenly mixed with inactive ports on the SUT; block designation (i.e., four active, then four inactive ports) is *not* acceptable.

Note: Packet loss during this run (if seen) invalidates the measurement and resets testing to Test Procedure 1 to provide a better L_{max} estimate. In this run, the SUT is *not* required to be able to return to higher performance levels instantly and *might* require a

⁸ The measurement interval represents a compromise between accuracy and speed. Shorter time intervals (i.e., 300 seconds) might 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.

⁹ Please refer to Appendix A for measurement conditions and qualifications.

configuration change and recovery period to return to full capacity

Average recorded energy consumption R_{50} for the same 1200 second interval should be read from the SUT at the end of this run with the appropriate means (SNMP, DMI, XML agent, command line interface, etc.).

Recommended: Average recorded system utilization U_{50} for the same 1200 second interval should be read from the SUT at the end of this run with the appropriate means (SNMP, DMI, XML agent, command line interface, etc.).

Step 4 (25% load) – this step can be also combined with “Test procedure 2, step 2”

This test run further reduces the load to L_{25} ($L_{25} = 0.25 \times L_{max}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption P_{25} calculated. Load reduction is achieved by either reducing traffic rate across all ports or by *sending traffic at full rate to every fourth port of the SUT*. Unused router tester ports can be left idle or shut down at vendor discretion. Active ports should be evenly mixed with inactive ports on SUT; block designation (i.e., 4 active, then 12 inactive ports) is *not* acceptable.

Packet loss during this run (if seen) invalidates the measurement and resets testing to Test Procedure 1 to provide a better L_{max} estimate. In this test, the SUT is *not* required to be able to return to higher performance levels instantly and *might* require a configuration change and recovery period to return to full capacity.

Average recorded energy consumption R_{25} for the same 1200 second interval should be read from the SUT at the end of this run with the appropriate means (SNMP, DMI, XML agent, command line interface, etc.).

Recommended: Average recorded system utilization U_{25} for the same 1200 second interval should be read from the SUT at the end of this run with the appropriate means (SNMP, DMI, XML agent, command line interface, etc.).

Step 5 (one-tenth ports in use) – this step can be combined w/ “Test procedure 2, step 3”

This test run further reduces the load to L_{10} ($L_{10} = 0.1 \times L_{max}$) and runs for another 1200 seconds. Energy consumption is being measured for the entire period and average consumption P_{10} calculated. Load reduction is achieved by either reducing traffic rate across all ports or by *sending traffic at full rate to every tenth port of the SUT*. Unused router tester ports can be left idle or shut down at vendor discretion. Active ports should be evenly mixed with inactive ports on SUT; block designation (i.e., 10 active, and then 90 inactive ports) is *not* acceptable.

Packet loss during this run (if seen) invalidates the measurement and resets testing to Test Procedure 1 to provide a better L_{max} estimate. In this test, the SUT is *not* required to be able to return to higher performance levels instantly and *might* require a configuration change and recovery period to return to full capacity

Average recorded energy consumption R_{10} for the same 1200 second interval should be read from the SUT at the end of this run with the appropriate means (SNMP, DMI, XML agent, command line interface, etc.).

Recommended: Average recorded system utilization U_{10} for the same 1200 second interval should be read from the SUT at the end of this run with the appropriate means (SNMP, DMI, XML agent, command line interface, etc.).

Test Procedure 5 (Optional Functional Test)

Collateral energy management

Apart from energy utilization and management within network and telecom devices themselves, they can also have an impact on energy consumption in connected (cascaded) devices.

Examples of this functionality can include the:

1. Ability to control power states in Power over Ethernet (PoE) connected devices
2. Ability to monitor energy consumption in PoE connected devices
3. Ability to assist power states in LAN-connected devices (i.e., data center cluster reconfiguration based on wake-on-LAN signal propagation)
4. Ability to control IP and non-IP devices via remote control mechanisms (such as BBF TR-069), including home and office energy monitoring and management

This section is a stub for a functional test of collateral energy management capabilities.

Effective Throughput Calculation

To normalize SUT energy consumption to performance, the latter number needs to be represented as throughput T_f . Two methods are used to convert load L_{max} (as recorded during the tests) to effective full-duplex throughput T_f (expressed in Gbps).

In the first method, the traffic generator (router tester) reports L_{max} as a combination of egress 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 are added to compute the effective wire rate at which the SUT performed in the test.

Note: The idle timeouts inserted by the SUT to compensate for asymmetric test patterns are *not* accounted for.

Example 1

The SUT is an Ethernet switch that can drive ten 10GbE ports at 7,291,702 frames per second each with 64Byte Ethernet frames without loss. According to the tester, this corresponds to 7,291,702 pps egress rate per each port.

$$T_f = 10 \times 7,291,702 \times 8 \times (64 + 1 + 7 + 12) = 49.000237440 \text{ Gbps}$$

(accounting for Ethernet start of frame, preamble and minimum interpacket gap)

In the second method, the tester equipment itself can report L_{\max} as the highest achieved line utilization on a per-port basis (in a percentage). 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 ten 10GbE ports (LAN PHY) on the access side (toward customer premises equipment) and ten 10GbE ports on the network side (towards MPLS core network).

The SUT can forward the incoming L2 frames (256 bytes each) towards the MPLS core with egress interface utilization of 100%. However, because of the 1:1 matching of access and network sides, the access side can be utilized only at 99.22% to allow lossless application of the (minimally) 4-byte MPLS L2 VPN header required for packets on the network side. The same limitation is seen in the opposite direction where the incoming network-side packets can fill only the access-side interfaces at 99.22% after the headers are stripped. Note, that network side (with all minimally required headers) is considered to be loaded at 100 percent in the egress direction (MPLS VPN header is considered L2 in this example):

The data rate for 10GbE (IEEE 802.3ae) is 10,000 Mbps:

$$T_f = 10 \times 10,000 \times 1.0 + 10 \times 10,000 \times 0.9922 = 199,220 \text{ Mbps} = 199.22 \text{ Gbps}$$

Example 3

The SUT is an Ethernet switch with eight GbE ports that can operate at 100% line utilization when configured for 802.1q packet encapsulation (VLAN headers applied). The same switch can operate only at 90% line rate utilization when not configured for VLAN encapsulation.

The measurement results from the second case should be used, because equipment class 3 (Ethernet L2/L3 switch devices) does not require VLAN headers to be present and they are not considered to be a necessary overhead.

APPENDIX B: Measurement Conditions

A.1 Temperature

The equipment should be evaluated at an ambient temperature of $25^{\circ} \pm 3^{\circ}\text{C}$. 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 should be evaluated at a relative humidity of 30% to 75%.

A.3 Air pressure

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

A.4 DC voltage

The input to the SUT (all active feeds) should be at a nominal DC voltage $\pm 5\%$. This corresponds to -53 ± 2.65 V for typical telecom facilities.

A.5 AC voltage and frequency

The input to the SUT (all active feeds) should be the specified voltage $\pm 1\%$ and the specified frequency $\pm 1\%$

A.6 Metrology requirements

Every active power feed should have the power (amp) meter installed inline with desired accuracy no less than $\pm 1\%$ of the actual power consumption. This should include correction for power factor (PF) on AC feeds.

A.7 Sampling frequency

All energy consumption calculations are based on averaging multiple readings over the course of measurements. Power meters should be able to produce no less than 100 evenly-spaced readings in every full test cycle duration.

APPENDIX C: Proposed Product Classes and 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.

Topologies in Use

T1 (full mesh)

Symmetric full-mesh topology (same-bandwidth traffic stream from every port to every port).

This topology is applicable to equipment with equal port roles and minimal internal oversubscription, such as core routers and carrier Ethernet switches (Fig. 1).

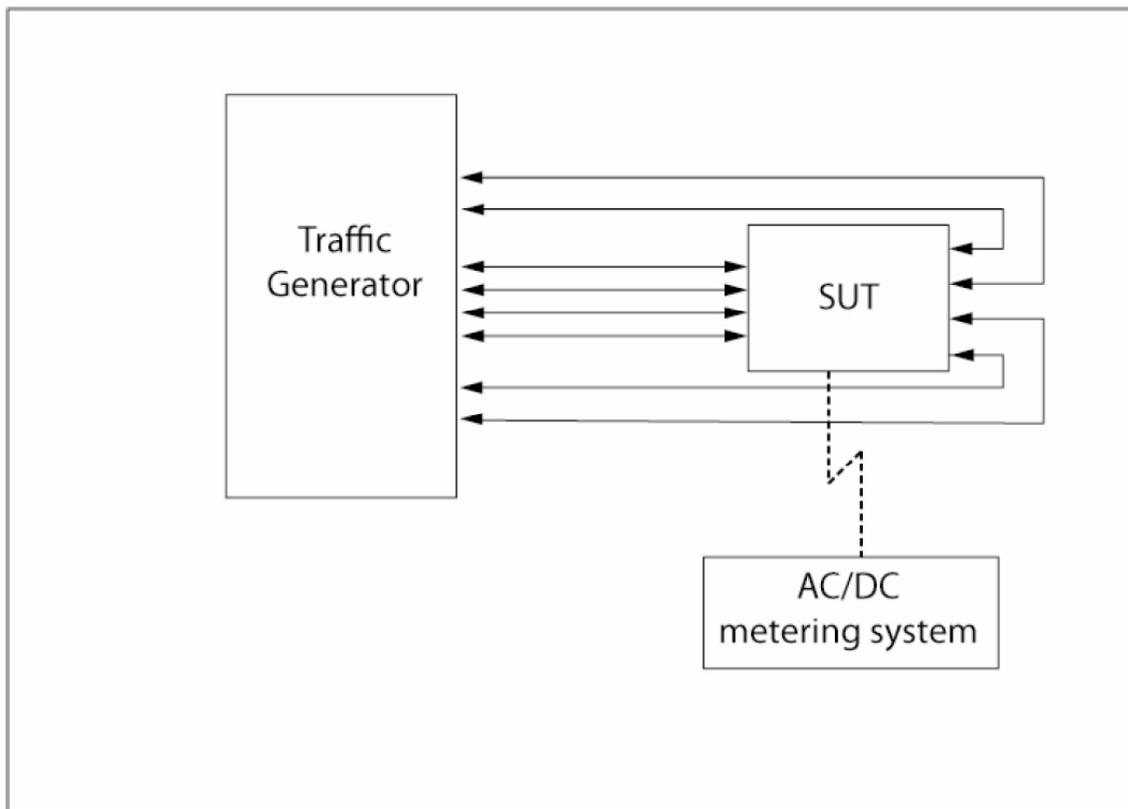


Figure 1. Full-mesh connection topology

T2 (dual-group partial mesh)

The ports on the SUT can be organized into network and access groups according to vendor discretion. Every network side port should be configured to send to every access side port and vice versa. All traffic streams on each side (network or access) are required to be of the same capacity. No traffic is allowed between ports of the same group (Fig. 2).

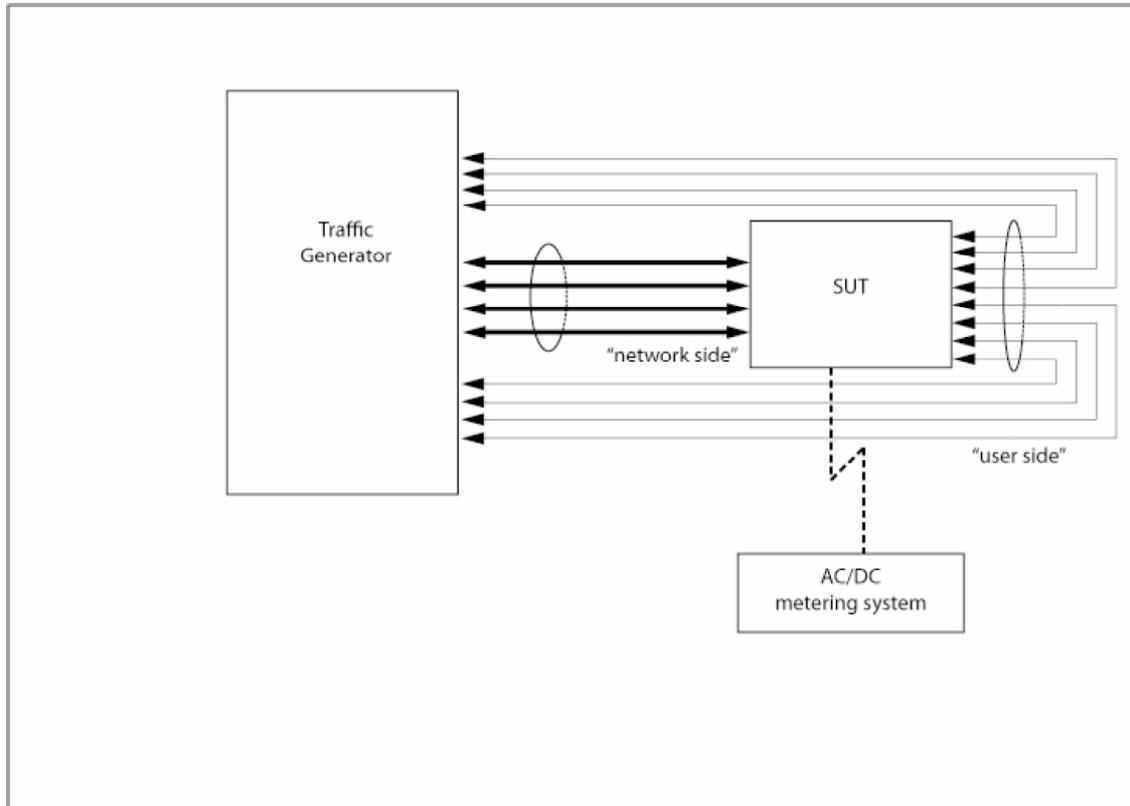


Figure 2. Dual-group partial mesh topology

This topology is applicable to various aggregation and edge transport devices where traffic from one side (user or access) is groomed or tunneled over to the other side (network). Note, that in this topology, the sum of face values of the “user side” ports can be significantly higher than sum of “network side” ports due to oversubscription.

However, for purposes of throughput calculation, only sustained egress data rates are accounted for.

Class 1—Routers

C1.1 Core Routers

Description: Core routing platforms are systems with terabit (half-duplex) or higher capacity and comprehensive system redundancy. 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: 500 Gbps or better full-duplex capacity, IP and/or MPLS support

Test application: IPv4, IPv6 or MPLS forwarding at discretion of the vendor; L3 packet size (MPLS considered L3): 340 Bytes (simple iMix average) or smaller; forwarding over any types of forwarding entries (static, connected, IGP, EGP)—no less than 64 active routes

Interface types: SONET, 10GbE or 100GbE 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

Test topology: T1

C1.2 Carrier Edge Routers/Ethernet Service Routers

Description: Carrier-grade edge routing platforms

Qualification: (IP and/or MPLS) VPN capability

Test applications: IPv4/IPv6 VPN, Ethernet pseudowire (PWE) or VPLS forwarding at discretion of the vendor; payload packet size: 340 Bytes (simple iMix average) or smaller; (IPv4, IPv6 or Ethernet frames as delivered to/from access side are considered to be payload); forwarding over any types of forwarding entries across no less than 128 VPN instances, with no less than 2048 separate VPN destinations active in total (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.

Test topology: T2. For purposes of VPN forwarding test, every access-side port should belong to all VPN instances

***Example 1** An Ethernet Services Router has twenty (20) GbE ports on the access side and two (2) 10GbE ports on the network side. Every access port on the SUT is divided into 128 VLANs, each VLAN belonging to one of the 128 VPLS instances configured. The router tester simulates no less than one MAC address (simulated host) per each local (access side) and remote (network side) Ethernet segments. Router tester simulates one remote PE (with 128 VPLS segments) connected to each of the two network-side ports of SUT*

In the traffic profile, every simulated host on “access side” sends two (2) equal-rate streams towards “network side” (one per remote PE). Every simulated host on the “network side” sends twenty (20) equal-rate streams towards “access side”, one per every “access side” port

C1.3 Multipurpose Routers

Description: Routing platforms of variable purposes (enterprise, edge, etc.)

Qualification: Any type of L3 packet forwarding

Test applications: IPv4 or IPv6 forwarding at vendor discretion. L3 packet size: 340 Bytes (simple iMix average) or smaller; forwarding over any types of forwarding entries, no less than 16 K active routes

Interface types: Electrical or optical at vendor discretion

Redundancy: For the purposes of testing, redundant components might *not* be present

Topology: T1 or T2 at vendor discretion

Class 2—WAN/Broadband Aggregation Device

C2.1 BRAS Devices

Description: Legacy broadband aggregation devices

Qualification: Point-to-Point Protocol over Ethernet (PPPoE), Point-to-Point Protocol over ATM (PPPoA) or Point-to-Point Protocol (PPP) session termination. Quality of service (QoS) provisioned at subscriber level

Test applications: PPPoE, PPPoA, PPP forwarding at discretion of the vendor; L3 packet size: 340 Bytes (simple iMix average) or smaller; forwarding over any types of per-subscriber entries, no less than 64,000 subscribers with no less than 4 queues assigned to each

Interface types: Short reach (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.

Topology: T2

C2.2 BNG/Common Edge devices

Description: Broadband aggregation devices, Ethernet-oriented

Qualification: PPPoE, PPP or IP DHCP session termination, per subscriber QoS

Test applications: IP DHCP, PPPoE, PPP forwarding at discretion of the vendor; L3 packet size: 340 Bytes (simple iMix average) or smaller; (measured as IPv4 or IPv6 payload to/from access side); forwarding over any types of per-subscriber entries, no less than 64,000 subscribers with no less than 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

Topology: T2

Class 3—Ethernet L2/L3 Switches

C3.1 Core/Datacenter Ethernet Switching Platforms

Description: High-performance Ethernet switching platforms

Qualification: L2 (Ethernet or MPLS) forwarding, L3 (IPv4 or IPv6 forwarding).

Test application: L2 or L3 forwarding at vendor discretion; payload packet size: 340 Bytes (simple iMix average) or smaller L2 or L3 frames; forwarding over any types of forwarding entries and encapsulation types

Interface types: SR optical (10/100/1000MbE or 10/100GbE) 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

Topology: T1

C3.2 Desktop/Aggregation Ethernet Platforms

Description: Ethernet switching platforms

Qualification: L2 (Ethernet) forwarding, IPv4, or IPv6 forwarding

Test application: Ethernet, IPv4/v6 forwarding at vendor discretion; payload packet size: 340 Bytes (simple iMix average) or smaller L2 or L3 frames; forwarding over any types of forwarding entries and encapsulation types

Interface types: Copper or SR optical (10/100/1000/10GbE) at vendor discretion

Redundancy: For the purposes of testing, redundant components *can be* removed.

Topology: T2

Class 4—Experimental

Placeholder for any equipment type are not assigned to a particular class.

Results in this category can be reported along with precise description of the test methodology, topology, traffic profile and protocol configuration

Class 5—Security Appliances (DPI, Firewalls, VPN Gateways)

Description: Security platforms of variable purposes (IPsec VPN, HTTPS, deep packet inspection (DPI), intrusion detection service (IDS), etc.)

Qualification: L3 forwarding, security features

Test application: IPsec or HTTPS, minimum number of firewall or DPI forwarding rules at vendor discretion; 340 Bytes (simple iMix average) or smaller L3 packets

Interface types: At vendor discretion

Redundancy: For the purposes of testing, redundant components *can be* removed.

Topology: TBD

Class 6—Application Gateways (Layers 5 Through 7 Accelerators, Load Balancers)

Description: Application platforms of variable purposes (server load balancer [SLB], accelerators, compressors)

Qualification: Application-specific features

Test application: User traffic at vendor discretion (need more qualification for setup); 340 Bytes (simple iMix average) or smaller L3 packets

Interface types: At vendor discretion

Redundancy: For the purposes of testing, redundant components *can be* removed

Topology: TBD

References

[ATIS-0600015.XX.2009] *Energy Efficiency for Telecommunication Equipment: Methodology for Measurement and Reporting for Router and Ethernet Switch Products.*

[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.

[Globecom 2009]. Kharitonov, Daniel. *Time-Domain Approach to Energy Efficiency: High-Performance Network Elements in Design.* Proceedings of IEEE Globecom 2009.

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

[JNPR TSBG] “Two Steps Beyond Greenwashing.”
© 2008 Juniper Networks, Inc.

[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]. Ceuppens, Luc, Kharitonov, Daniel, and Sardella, Alan. *Power Saving Strategies and Technologies in Network Equipment Opportunities and Challenges, Risk and Rewards.* Proceedings of IEEE Symposium on Applications and the Internet.

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

THIS PAGE IS INTENTIONALLY LEFT BLANK