

Benchmarking Methodology Working Group
Internet-Draft
Intended status: Standards Track
Expires: 16 August 2026

C. Pignataro
NC State University
R. Jacob
ETH Z^{urich}
G. Fioccola
Q. Wu
G. Chen
Huawei
S. Prabhu
Nokia
12 February 2026

Characterization and Benchmarking Methodology for Power in Networking
Devices
draft-ietf-bmwg-powerbench-01

Abstract

This document defines a standard mechanism to measure, report, and compare power usage of different networking devices under different network configurations and conditions.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 16 August 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document.

Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Aim and Scope	3
3. Replicability and Comparability	4
4. Terminology	4
4.1. Total Weighted Capacity of the interfaces	4
4.2. Total Weighted Power	5
4.3. Energy Efficiency Ratio	6
4.4. Total Power	7
5. Energy Consumption Benchmarking	8
6. Test Methodology	8
6.1. Test Setup	8
6.2. Traffic and Device Characterization	9
7. Reporting Format	10
8. Benchmarking Tests	11
8.1. Throughput	11
8.2. Base Power	12
8.3. Idle Power	12
8.4. Idle+ Power	13
8.5. Power with Traffic Load	13
8.6. Energy Efficiency Ratio	14
9. Security Considerations	14
10. IANA Considerations	15
11. Acknowledgements	15
12. References	15
12.1. Normative References	15
12.2. Informative References	15
Authors' Addresses	16

1. Introduction

Energy efficiency is becoming increasingly important in the operation of network infrastructure. Network devices are typically always on, but in some cases, they run at very low average utilization rates. Both network utilization and energy consumption of these devices can be improved, and that starts with a normalized characterization [RFC7460].

The benchmarking methodology defined here will help operators to get a more accurate idea of the power drawn by their network and will also help vendors to test the energy efficiency of their devices [RFC6988].

There is no standard mechanism to benchmark the power utilization of networking devices like routers or switches.
[I-D.manral-bmwg-power-usage] started to analyze the issue.

This document focuses on the mechanism to correctly characterize and benchmark the energy consumption of networking devices to better estimate and compare their power usage in order to assess the performance over a set of well-defined scenario.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Aim and Scope

Benchmarking can be understood to serve three related but different objectives:

Assessing "which system performs best" over a set of well-defined scenarios.

Measuring the contribution of sub-systems to the overall system's performance (also known as "micro-benchmark").

Evaluating the evolution of a system's energy consumption profile against its own historical baseline to assess the impact of software updates, feature activation, or operational aging (also known as "longitudinal" or "self-referential" benchmarking).

Achieving any of these objectives requires a well-defined set of principles prescribing what must be measured, how must be measured, and which results must be reported. Providing those principles is the objective of this draft. These are simply called "the benchmark" in the rest of this draft.

The benchmark aims to compare the energy efficiency for individual devices (routers and switches belonging to similar device classes). In addition, it aims to showcase the effectiveness of various energy optimization techniques for a given device and load type, with the objective of fostering improvements in the energy efficiency of future generations of devices.

3. Replicability and Comparability

Replicability is defined as achieving the same results with newly collected data. Formally, it is a pre-requisite for benchmarking. Benchmark results are meant to be compared, and this comparison is not sound if the individual results are not replicable.

As discussed later in this draft, replicability in power measurements is complex as power is affected by a wide range of parameters, some of which are hard to control e.g., the room temperature.

Striving for "perfect" replicability would lead to prescribe extremely precisely all the power-impacting factors in the test setup. We argue that this is unrealistic and counter-productive. An overly prescriptive benchmark becomes more complicated to perform. Furthermore, results would then be comparable only across benchmark results obtained under the exact same test conditions, which becomes increasingly less likely as we prescribe more and more.

Instead, the benchmark described in this draft proposes to report on a number of power-impacting factors, but does not enforce specific values or settings for those. The aim is to make the benchmark easier to perform. The comparison between benchmark results MAY be somewhat less accurate or fair than with a more prescriptive benchmark, but the hope is to have many more comparison points available, which would ultimately provide a more robust image of the devices power demands and their evolution over time.

In short: this draft argues it is better to have many benchmark results with a higher uncertainty than a few very precise but hardly comparable ones.

4. Terminology

4.1. Total Weighted Capacity of the interfaces

The total weighted capacity of the interfaces (T) is the weighted sum of all interface throughputs.

Definition:

$$T = B1*T1 + \dots + Bi*Ti + \dots + Bm*Tm$$

Discussion:

Ti is the total capacity of the interfaces for a fixed configuration model and traffic load (the sum of the interface bandwidths)

Bi is the weighted multiplier for different traffic levels (note that $B1 + \dots + Bj + \dots + Bm = 1$, weight multipliers MAY be specified for router, switch differently, 3 typical weighted multipliers are 0.1,0.8,0.1)

m is the number of traffic load levels (if it is considered 100%, 30%, 0%; $m = 3$) Note that traffic load levels MAY be specified differently for router and switch, e.g., traffic level 100%,10%,0% for access router, traffic level 100%,30%,0% for core router and data center switch.

Measurement units:

Gbps.

Issues

The traffic loads and the weighted multipliers need to be clearly established a priori.

It is unclear if the definition of the Ti is/should be linked to the traffic load levels. For a given port configuration (which may result in 50% of the total capacity a device can provide), one may be interested in traffic load of e.g., 5% or 10% or the total capacity (not only 50%).

See Also:

[ETSI-ES-203-136],[ITUT-L.1310] ,[ATIS-0600015.03.2013].

4.2. Total Weighted Power

The total weighted power (P) is the weighted sum of all power calculated for different traffic loads.

Definition:

$$P = B1*P1 + \dots + Bi*Pi + \dots + Bm*Pm$$

Discussion:

P_i is the Power of the equipment in each traffic load level (e.g. 100%, 30%, 0%)

B_i is the weighted multiplier for different traffic levels (note that $B_1 + \dots + B_j + \dots + B_m = 1$)

m is the number of traffic load levels (if it is considered 100%, 30%, 0%; $m = 3$)

Measurement units:

Watt.

Issues:

The traffic loads and the weighted multipliers need to be clearly established a priori.

Importantly, the traffic MUST be forwarded of the correct port! It would be easy to cut power down by dropping all traffic, and, naturally, we do not want that. A tolerance on packet loss and/or forwarding error must be specified somehow. That tolerance could be zero for some benchmark problems (e.g., Non-Drop Rate (NDR) estimation), and non-zero for others. Tolerating some errors may be interesting to navigate the design space of energy saving techniques, such as approximate computing/routing. According to measurement procedure in section 6.5 of [ATIS-0600015.03.2013], the Equipment Under Test (EUT) should be able to return to full NDR load. Failure to do so disqualifies the test results.

See Also:

[ETSI-ES-203-136], [ITUT-L.1310] , [ATIS-0600015.03.2013].

4.3. Energy Efficiency Ratio

Energy Efficiency Ratio (EER) is defined as the throughput forwarded by 1 watt and it is introduced in [ETSI-ES-203-136]. A higher EER corresponds to a better energy efficiency.

Definition:

$$EER = T/P$$

Discussion:

T is the total weighted sum of all interface throughputs

P is the weighted power for different traffic loads

Measurement units:

Gbps/Watt.

Issues:

The traffic loads and the weighted multipliers need to be clearly established a priori.

Optionally the total capacity of the interfaces (T) can be used in replacement of the total weighted capacity of the interfaces. The former is the sum of all interface throughputs and is not linked to traffic load levels. This result EER is a scaled value proportional to the EER using the total weighted capacity of the interfaces (T)

See Also:

[ETSI-ES-203-136],[ITUT-L.1310] ,[ATIS-0600015.03.2013].

4.4. Total Power

The total power (P_{tot}) is the power of the entire equipment, measured as the sum the power drawn by all of the equipment's power supply units.

Definition:

$$P_{tot} = P_{u1} + \dots + P_{ui} + \dots + P_{un}$$

Discussion:

P_{ui} is the power that is drawn by one power supply unit of the equipment

n is the number of power supply units

Measurement units:

Watt.

Issues:

The total power depends on many different factors, including the running configuration, the number and type of transceiver connected, the forward traffic volume and pattern, the version of

the operating system, the room temperature and humidity/other environmental dimensions, the aging of parts, etc. This metric does not allow to compare two equipments against each other, but it MAY be enough to assess the effect of a change on the same equipment; e.g., for optimizing the power draw by changing the running configuration.

Importantly, the traffic MUST be forwarded of the correct port! It would be easy to cut power down by dropping all traffic, and we of course do not want that. A tolerance on packet loss and/or forwarding error MUST be specified somehow. That tolerance could be zero for some benchmark problems, and non-zero for others. Tolerating some errors MAY be interesting to navigate the design space of energy saving techniques, such as approximate computing/routing.

5. Energy Consumption Benchmarking

The maximum power drawn by a device does not accurately reflect the power under a normal workload. Indeed, the energy consumption of a networking device depends on its configuration, connected transceivers, and traffic load. Relying merely on the maximum rated power can overestimate the total energy of the networking devices.

A network device consists of many components, each of which draws power (for example, it is possible to mention the power draw of the CPU, data forwarding ASIC, memory, fan, etc.). Therefore, it is important to formulate a consistent benchmarking method for network devices and consider the workload variation and test conditions.

Enforcing controlled conditions on test conditions (e.g., Temperature) is important for test procedure to make sure test conditions repeatable [RFC6985]. The measurement condition reported in [ATIS-0600015.2009] and [ITUT-L.1310] SHOULD be applied, e.g., the power measurements SHALL be performed in a laboratory environment under specific range of temperature, humidity and atmosphere pressure.

6. Test Methodology

6.1. Test Setup

The test setup in general is compliant with [RFC2544]. The Device Under Test (DUT) is connected to a Tester and a Power Meter. The Power Meter allows measurement of the device's energy consumption and can be used to measure power under various configurations and conditions. Tests MUST be done by running one or several of the predefined traffic traces, which are crafted to test different power-

intensive tasks related to packet processing. The Tester is also a traffic generator that enables changing traffic conditions. It is OPTIONAL to choose a non-equal proportion for upstream and downstream traffic.




Figure 1: Test Setup

It is worth mentioning that the DUT also dissipates significant heat. A part of the power is used for actual work while the rest is dissipated as heat. This heating can lead to more power drawn by fans/compressor for cooling the devices. The benchmarking methodology does not measure the power drawn by external cooling infrastructure. The Power Meter only measures the internal energy consumption of the device. Anyway, the device's temperature change MUST be known. It is useful to know whether device's heat management plays a role in the observed differences in energy efficiency and can be correlated to the amount of external power drawn to cool the device.

6.2. Traffic and Device Characterization

The traffic load supported by a device affects its energy consumption. Therefore, the benchmark MUST include different traffic loads.

The traffic load MUST specify packet sizes, packet rates, and inter-packet delays, as all MAY affect the energy consumption of network devices [ATIS-0600015.2009]. To enable replicable and comparable results, the benchmark can specify a set of well-defined traffic traces that MUST be used.

There are different interface types on a network device and the power usage also depends on the kind of connector/transceiver used. The interface type used MUST also be specified.

Power measurements SHOULD be performed only after the DUT and the applied traffic condition have reached a stable operating state. Stability in this context refers to the absence of transient effects associated with traffic changes, configuration updates, or device thermal and cooling behavior.

The stabilization interval, defined as the time between applying a traffic load or configuration and the start of power measurement, MUST be reported.

The measurement interval, defined as the averaging window over which input power is recorded, and the averaging method used MUST also be reported.

This methodology does not mandate specific durations for stabilization or measurement intervals, as these MAY depend on device class, implementation, and test environment; however, reporting these intervals is required to support comparability of results.

7. Reporting Format

The benchmark focuses on data that is either controllable (e.g., the number of active ports) or that can be externally measured (e.g., the total power). Factors that are not measurable externally (e.g., CPU load, PSU efficiency) are intentionally left out.

Network Device Hardware and Software versions:

For the benchmarking tests, it MUST be specified.

Number and type of line cards:

For each test the total number of line cards and their types can be varied and MUST be specified.

Number of enabled ports:

For each test the number of enabled and disabled ports MUST be specified.

Number of active ports:

For each test the number of active and inactive ports MUST be specified.

Port settings and interface types:

For each test the port configuration and settings need to be specified.

Port Utilization:

For each test the port utilization of each port MUST be specified. The actual traffic load can use the information defined in [RFC2544]. To characterize power consumption under realistic traffic conditions, the actual traffic load can use variable packet size distributions as specified in [RFC6985].

Traffic trace:

For each test, the traffic trace used (amongst those prescribed by the benchmark) MUST be specified.

Power measurement:

For each test it MUST be specified. All power measurements are done in Watts.

Stabilization and measurement intervals:

For each power measurement, the stabilization interval and the measurement interval (averaging window) MUST be reported. The stabilization interval is the time between applying a traffic load or configuration and the start of power measurement. The measurement interval and the averaging method used to obtain the reported power value MUST also be specified. If these intervals differ by load level, the values per load level MUST be reported.

8. Benchmarking Tests

8.1. Throughput

Objective:

To determine the DUT throughput according to [RFC2544]. And to characterize throughput under realistic traffic patterns using variable packet size distributions as specified in [RFC6985].

Procedure:

The test is done using a multi-port setup as specified in Section 16 and Section 26.1 of [RFC2544]. To assess throughput under realistic Internet traffic conditions, the test is performed using the IMIX Genome packet size distributions specified in [RFC6985].

Reporting format:

The results of the throughput SHOULD be reported according to Section 7.

8.2. Base Power

Objective:

To determine the base power drawn by the network device in its factory settings.

Procedure:

The measurement is done with the device in its factory settings, after it finished booting, and without any transceiver plugged in.

Reporting format:

The results of the power measurement SHOULD be reported according to Section 7.

Note:

This measurement is useful to assess the energy efficiency of default settings.

8.3. Idle Power

Objective:

To determine the power drawn by the network device in normal operation but without forwarding traffic.

Procedure:

The measurement is done with the device fully configured to forward traffic but without any traffic actually present. All interfaces MUST be up.

Control-plane, management-plane, and background system functions MAY remain active during this measurement unless explicitly disabled for the purpose of the test. Examples include routing protocol adjacencies, control signaling, telemetry processes, and thermal management mechanisms. The operational state of these functions during the measurement MUST be reported, as they can influence the observed power consumption.

Reporting format:

The results of the power measurement SHOULD be reported according to Section 7.

Note:

This measurement is useful to assess the energy used to activate the internal components used by the device to forward traffic. It also captures the efficiency of the device at activating some "low-power mode" when there is no traffic to forward.

8.4. Idle+ Power**Objective:**

To determine the power drawn by the network device in normal operation with very small but non-zero traffic to forward.

Procedure:

The measurement is done with the device fully configured and the "minimum" traffic trace.

Control-plane, management-plane, and background system functions MAY remain active during this measurement unless explicitly disabled for the purpose of the test. Examples include routing protocol adjacencies, control signaling, telemetry processes, and thermal management mechanisms. The operational state of these functions during the measurement MUST be reported, as they can influence the observed power consumption.

Reporting format:

The results of the power measurement SHOULD be reported according to Section 7.

Note:

The "minimum" traffic trace creates a bidirectional flow of 1 pps on all active interfaces. By comparison with the "Idle Power" measurement, this measurement captures the power cost of taking the device out of its "low-power mode."

8.5. Power with Traffic Load**Objective:**

To determine the power drawn by a device. The dynamic power, which is added to the idle+ power, SHOULD be proportional to its traffic load.

Procedure:

A specific number of packets at a specific rate is sent to specific ports/linecards of the DUT. All DUT ports must operate under a specific traffic load, which is a percentage of the maximum throughput.

Power SHOULD be recorded only after the DUT and the offered load have reached a stable operating state. The stabilization interval and the measurement interval MUST be reported as described in Section 7.

Reporting format:

The results of the power measurement SHOULD be reported according to Section 7.

8.6. Energy Efficiency Ratio

Objective:

To determine the energy efficiency of the DUT.

Procedure:

Collect the data for all the traffic loads and apply the formula of Section 4. For example, with all DUT ports operating stably under a percentage of the maximum throughput (e.g. 100%, 30%, 0%), record the average input power and calculate the total weighted power P and then the EER.

The stabilization interval and measurement interval used to obtain the recorded average input power MUST be reported as described in Section 7.

Reporting format:

The results of the energy efficiency ratio SHOULD be reported according to Section 7.

9. Security Considerations

The benchmarking characterization described in this document is constrained to a controlled environment (as a laboratory) and includes controlled stimuli. The network under benchmarking MUST NOT be connected to production networks.

Beyond these, there are no specific security considerations within the scope of this document.

10. IANA Considerations

This document has no IANA actions.

11. Acknowledgements

We wish to thank the authors of [I-D.manral-bmwg-power-usage] for their analysis and start on this topic.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7460] Chandramouli, M., Claise, B., Schoening, B., Quittek, J., and T. Dietz, "Monitoring and Control MIB for Power and Energy", RFC 7460, DOI 10.17487/RFC7460, March 2015, <<https://www.rfc-editor.org/info/rfc7460>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

12.2. Informative References

- [ATIS-0600015.03.2013] ATIS, "ATIS-0600015.03.2013: Energy Efficiency for Telecommunication Equipment: Methodology for Measurement and Reporting for Router and Ethernet Switch Products", 2013.
- [ATIS-0600015.2009] ATIS, "ATIS-0600015.2009: Energy Efficiency for Telecommunication Equipment: Methodology for Measurement and Reporting - General Requirements", 2009.
- [ETSI-ES-203-136] ETSI, "ETSI ES 203 136: Environmental Engineering (EE); Measurement methods for energy efficiency of router and switch equipment", 2017, <https://www.etsi.org/deliver/etsi_es/203100_203199/203136/01.02.00_50/es_203136v010200m.pdf>.

- [I-D.manral-bmwg-power-usage]
Manral, V., Sharma, P., Banerjee, S., and Y. Ping,
"Benchmarking Power usage of networking devices", Work in
Progress, Internet-Draft, draft-manral-bmwg-power-usage-
04, 12 March 2013, <[https://datatracker.ietf.org/doc/html/
draft-manral-bmwg-power-usage-04](https://datatracker.ietf.org/doc/html/draft-manral-bmwg-power-usage-04)>.
- [ITUT-L.1310]
ITU-T, "L.1310 : Energy efficiency metrics and measurement
methods for telecommunication equipment", 2020,
<<https://www.itu.int/rec/T-REC-L.1310/en>>.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for
Network Interconnect Devices", RFC 2544,
DOI 10.17487/RFC2544, March 1999,
<<https://www.rfc-editor.org/info/rfc2544>>.
- [RFC6985] Morton, A., "IMIX Genome: Specification of Variable Packet
Sizes for Additional Testing", RFC 6985,
DOI 10.17487/RFC6985, July 2013,
<<https://www.rfc-editor.org/info/rfc6985>>.
- [RFC6988] Quittek, J., Ed., Chandramouli, M., Winter, R., Dietz, T.,
and B. Claise, "Requirements for Energy Management",
RFC 6988, DOI 10.17487/RFC6988, September 2013,
<<https://www.rfc-editor.org/info/rfc6988>>.

Authors' Addresses

Carlos Pignataro
North Carolina State University
United States of America
Email: cpignata@gmail.com, cmpignat@ncsu.edu

Romain Jacob
ETH Zürich
Switzerland
Email: jacobr@ethz.ch

Giuseppe Fioccola
Huawei
Italy
Email: giuseppe.fioccola@huawei.com

Qin Wu
Huawei
China
Email: bill.wu@huawei.com

Gen Chen
Huawei
China
Email: chengen@huawei.com

Shailesh Prabhu
Nokia
India
Email: shailesh.prabhu@nokia.com