
**Information technology —
Information technology sustainability
— Energy efficient computing
models —**

**Part 1:
Guidelines for energy effectiveness
evaluation**

*Technologies de l'information — Disponibilité des technologies de
l'information — Modèles informatisés à efficacité énergétique —*

*Partie 1: Lignes directrices pour l'évaluation de l'efficacité
énergétique*

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 30132-1:2016



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2016, Published in Switzerland

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Ch. de Blandonnet 8 • CP 401
CH-1214 Vernier, Geneva, Switzerland
Tel. +41 22 749 01 11
Fax +41 22 749 09 47
copyright@iso.org
www.iso.org

Contents

	Page
Foreword.....	iv
Introduction.....	v
1 Scope.....	1
2 Normative references.....	1
3 Terms and definitions.....	1
4 Abbreviated terms.....	1
5 Reference computing model for end-to-end energy effectiveness evaluation.....	2
5.1 Overview of computing models.....	2
5.2 Reference computing model and energy effectiveness evaluation.....	4
6 Holistic framework for evaluating the applicability of energy effectiveness improving technologies.....	6
6.1 Motivation.....	6
6.2 Overview of holistic framework.....	7
6.3 Considerations for evaluating the applicability.....	7
6.4 Examples of energy effectiveness evaluation.....	7
7 Guidelines for determining the energy effectiveness of a computing model.....	11
Annex A (informative) State of affairs for improving energy effectiveness of computing systems.....	12
Annex B (informative) Survey for calculating power consumption of the components in computing models.....	19
Bibliography.....	21

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: [Foreword - Supplementary information](#)

The committee responsible for this document is ISO/IEC JTC 1, *Information technology*, Subcommittee SC 39, *Sustainability for and by Information Technology*.

A list of all parts of the ISO/IEC 30132 series can be found on the ISO website.

Introduction

The world is experiencing explosive growth of data from mobile client devices, cloud services, social networks, online television, the Internet of things, big data and from traditional enterprise computing. The growth of data has been accompanied by a growth in the energy usage and carbon footprint of IT along with increased costs. Much research has been performed regarding energy management for the last two decades, most focusing on the evaluating and improving energy efficiency of individual components or systems such as processors, memory, wireless networks base stations, laptops, supercomputers, data centres, handheld devices and so on. However, several disparate systems, or systems of systems, collectively use energy to accomplish a given task and satisfy service-level expectations. Consider, for example, someone who takes a photo with a smartphone and posts it to a social network for their friends to view. Taking and transmitting the photo consumes energy from the smartphone while the data transfer, processing and storage consumes energy too. Likewise, when friends view the photo, that activity will consume additional energy. To improve energy effectiveness, it is necessary to consider the end-to-end energy use of a task or service involving multiple systems.

The ISO/IEC 30132 series provides guidelines for the end-to-end evaluation of energy effectiveness of a reference computing model and suggestions for determining the energy effectiveness of a computing model. This document comprises guidelines for energy effectiveness evaluation, including a reference computing model that includes end-to-end data transfer, processing and storage.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 30132-1:2016

Information technology — Information technology sustainability — Energy efficient computing models —

Part 1: Guidelines for energy effectiveness evaluation

1 Scope

This document establishes guidelines for improving the energy effectiveness for computing models. Specifically, this document provides

- a reference computing model for evaluating end-to-end energy effectiveness,
- a holistic framework for evaluating the applicability of energy effectiveness improving technologies, and
- guidelines for evaluating energy effectiveness.

2 Normative references

There are no normative references in this document.

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 13273-1 and the following apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <https://www.iso.org/obp/>

3.1

energy effectiveness

end-to-end total amount of data transferred, processed and stored per unit energy of a computing model

4 Abbreviated terms

ARP	address resolution protocol
BNG	broadband network gateway
DHCP	dynamic host configuration protocol
DSL	digital subscriber loop
DSLAM	DSL access multiplexer
FTTN	fibre-to-the-node
FTTP	fibres-to-the-premises

HSPA	high speed packet access
ICMP	internet control message protocol
ICT	information and communication technology
ISP	internet service provider
NIC	network interface card
OLT	optical line terminal
ONU	optical network unit
PON	passive optical network
PtP	point-to-point
QoS	quality of service
UMTS	universal mobile telecommunications system
WDM	wavelength-division multiplexing

5 Reference computing model for end-to-end energy effectiveness evaluation

5.1 Overview of computing models

This subclause provides a survey of trends for various computing models.

In the traditional client-server computing model, clients are connected to servers via networks such as the Internet. In this paradigm, a server is a computer system that selectively shares its resources, whereas a client is a computer that initiates requests to a server in order to use its resources.

However, the emergence of new computing models such as cloud computing and the Internet of things, along with new devices such as mobile phones, tablets, wearables and Internet connected sensors introduce new considerations for both computing models and determining energy effectiveness.

Additionally, data is now transported over traditional wired networks and over high- and low-speed wireless networks. Some client devices only support wireless networks.

Energy effectiveness has traditionally been viewed on a per device-category basis, but now it is important to look at energy effectiveness from an end-to-end perspective that includes all the devices, sub-systems and software that delivers a given set of functionalities (also known as a service).

New paradigms for the creation and use of data also affect energy effectiveness end-to-end. For example, many users have their data on multiple devices and synchronization between client devices and servers has become common practice. This means the same data may reside on multiple devices. On the other hand, some applications retrieve and display data on the client device only as needed. This scenario increases loads on networks, servers and storage, which may increase energy consumption.

The shift to mobile client devices and the shift to cloud computing, along with increases in the total number of connected devices have driven a dramatic increase in the number of data centres and faster, higher capacity networks with a corresponding increase in energy use, while client devices continue to improve their energy effectiveness. Customer expectations for highly available, responsive services may cause servers, storage and networking equipment to stay at high power states longer, possibly conflicting with power management schemes and energy effectiveness goals. New technologies such as push notifications also increase data and energy use across systems.

Therefore, energy effectiveness assessments should identify all of the energy consuming components in the computing model. The energy effectiveness of networks is calculated using manufacturers' data on equipment energy consumption for a range of typical types of equipment in networks. The manufacturers' data can include the amount of energy consumption in various states of equipment such as idle, active and fully utilized state. This approach enables an overall model of network power consumption to be constructed and provides a platform for predicting the growth in power consumption as the number of users and access rate per user increase[23].

[Figure 1](#) shows a high level representation of the network model of the Internet. In [Figure 1](#), the Internet is segmented into three major components: access network, metro network and core network with data centres. The model is an abstract representation of the Internet and, as such, does not include much of the fine detail of the Internet's true structure and topology. The model does account for the typical hop count for packets that traverse the Internet[24].

The refinement to include a more realistic representation of the Internet's topology is ongoing. The access network connects individual users to their local exchanges. Some of the typical access network technologies, such as digital subscriber loop (DSL) to deliver packets through fixed-line telephone service, fibres-to-the-premises (FTTP) installations to provide shared passive optical network (PON) or a point-to-point (PtP) Ethernet connection. In a PON, a single fibre from the network node feeds one or more clusters of users by using a passive optical splitter. An optical line terminal (OLT) is located at the local exchange to serve many access modems or optical network units (ONUs) located at each user. ONUs communicate with the OLT in a time division multiplexing, with the OLT assigning time slots to each ONU based on its relative demand. In a PtP access network, each ONU is directly connected to the local exchange with a dedicated fibre to the exchange. In areas where the copper pairs are in good condition, a fibre-to-the-node (FTTN) technology may be used. This technology uses a dedicated fibre from the local exchange to a DSL access multiplexer (DSLAM) located in a street cabinet close to a cluster of users. A high-speed copper pair technology, such as very-high-speed DSL, is used from the cabinet to the users. In areas where copper and fibre are not available or feasible, wireless can provide Internet access. Technologies for the wireless access include WiMAX, high speed packet access (HSPA), and universal mobile telecommunications system (UMTS). For wireless access, a wireless modem, located in the user, communicates with a local wireless base station which, in turn, is connected to the central office[23].

The central offices in a city are connected to each other and to other cities via the metro/edge network. This network also provides connection points for Internet service providers (ISPs). The metro and edge network serves as the interface between the access and core networks. The metro and edge network includes edge Ethernet switches, broadband network gateway (BNG) and provider edge routers. Edge Ethernet switches concentrate traffic from a large number of access nodes uplink to two or more BNG routers. The edge switch connects to two or more BNG routers to provide redundancy. The BNG routers perform access rate control, authentication and security services, and connect to multiple provider edge routers to increase reliability. The provider edge routers connect to the core of the network[23].

The core network comprises a small number of large routers in major population centres. These core routers perform all the necessary routing and also serve as the gateway to neighbouring core nodes. The core routers of any one network are often highly meshed, but only have few links to the networks of other providers. High-capacity wavelength-division multiplexed (WDM) fibre links interconnect these routers and connect to networks of other operators[23].




Figure 1 — High-level network structure for the Internet^[23]

5.2 Reference computing model and energy effectiveness evaluation

Since there are many components in modern computing models and each has unique energy effectiveness characteristics, it is difficult to calculate the energy consumption of services. Therefore, this document uses a simplified generic reference computing model consisting of client devices, network equipment and data centre equipment as shown in Figure 2. It is assumed that data centre has computing resources such as server and storage. This document considers the following assumptions for the simple evaluation of energy effectiveness of computing models.

- This document evaluates energy effectiveness of computing models from the view point of end-to-end data creation, processing, storage, consumption and sharing.
- The model also assumes a data-oriented scenario where data moves back and forth between client devices and data centres.

This document considers following components for evaluating end-to-end energy effectiveness while considering application's characteristics:

- client devices;
- network equipment;
- data centre equipment.

For survey for calculating power consumption of the components in computing models, see Annex B.




Figure 2 — Simplified generic reference computing model for energy effectiveness evaluation

From the perspective of computing the overall energy effectiveness, this document takes a very high-level macroscopic view. At this level, it is considered that the transport network consists of multiple wired and wireless network equipment. Client devices can take many forms including wearable technologies, sensors and controls. It is noted that this document considers network layer equipment in estimating the energy effectiveness of networks because the detail information about networking equipment below network layer is difficult to obtain due to security and confidentiality. Thus, internal network equipment are not taken into account. In case of data centre, border routers that connect data centre to the Internet are taken into account. The internal routers within a data centre are not considered. However, if the internal network information of clients, transport networks and data centre are available, those information may be utilized in order to estimate the energy effectiveness more accurately. Therefore, the end-to-end energy effectiveness evaluation can be performed by calculating the energy effectiveness of the following components:

- client devices;
- transport network: in [Figure 2](#), it is assumed that transport network includes access, metro/edge and core network;
- data centre equipment.

The individual architectures that make up the high level components in the model are much more complicated and the corresponding details and complexities are not within the scope of this document.

The energy consumption (watt-hours) of the end-to-end path in the reference computing model can be calculated as follows:

$$E_{\text{E2E}} = \sum_i E_{\text{client_device},i} + \sum_j E_{\text{net_equip},j} + \sum_k E_{\text{datacentre_equip},k}$$

The energy consumption measurement of components in the reference computing model can be performed as follows.

- Basic assumption: each component is equipped with monitoring function, which can monitor and measure throughput level of target equipment. When the monitoring value reaches a designated threshold value, the amount of consumed energy is measured.
- Client devices and data centre equipment: monitoring function can measure the amount of power consumption at the offered load to the equipment. Calculation of the energy consumption on client

device or data centre equipment can be performed by using measurement time, T , and power consumption upon offered load as follows:

$$E_{\text{equip}} = \sum_0^T P_{\text{equip}} (\text{offered_load})$$

- Network equipment: network throughput can be calculated as doing sum of traffic load entered into each network equipment. For this, the monitoring function can collect information about network throughput such as offered load to measure the power information of the equipment. When the network throughput reaches at target threshold value for power measurement, the monitoring function gathers the information upon traffic load on each node and calculates the energy consumption at the target traffic load. The total amount of power consumption on the network equipment can be calculated by using the information upon each traffic load. The consumed power on total network equipment can be the sum of the consumed power on each node. The power measurement can be done depending on different throughput states (e.g. 10 %, 20 %, etc.). When the value of traffic load is known, any network equipment can calculate power consumption value of each node according to traffic load. It is noted that if calculating power consumption of each network equipment, a practical way to calculate the network energy consumption may be to take each segment, measure the total energy consumed over T , divide that by the total number of bytes of data that traversed that segment over T and then multiply by the number of bytes of data for the service that traversed that network segment over T .

$$E_{\text{net_equip}} = \sum_0^T P_{\text{net_equip}} (\text{traffic_load})$$

Since the energy effectiveness evaluation of reference computing model only deals with end-to-end data scenario, the end-to-end energy effectiveness can be calculated as follows.

Comparing product metrics allow consumers, enterprises and carriers to add energy effectiveness to purchase criteria. A straightforward way to estimate the energy effectiveness of a network or telecom system is to normalize its energy consumption to the amount of transmitted data in the test.

Let B_{E2E} be the end-to-end total amount of data transferred, processed and stored during the measurement time, T , in the reference computing model and η_{E2E} be the overall energy effectiveness calculated. Therefore, the end-to-end network energy effectiveness is calculated by dividing the total amount of data, i.e. B_{E2E} by the total energy consumption of the computing model.

$$\eta_{\text{E2E}} = \frac{B_{\text{E2E}}}{E_{\text{E2E}}}$$

6 Holistic framework for evaluating the applicability of energy effectiveness improving technologies

6.1 Motivation

As the energy effectiveness of computing systems that use a large amount of electricity becomes an important issue to network providers and operators, there is an increasing demand for developing standards for evaluating and assessing the effectiveness and applicability of various energy effectiveness improving technologies. This clause defines effective and holistic framework for assessing the energy effectiveness improving technologies.

The rapid increase of energy consumption of ICT products accelerates the development of energy efficient technologies. However, since energy efficient technologies in different dimensions may cause unexpected side effects in total energy effectiveness of a computing system, it becomes necessary to develop effective methods for understanding the inter-relationship among multiple energy efficient technologies and evaluating energy effectiveness based on holistic view. This subclause also provides motivation for holistic framework.

For state of affairs for improving energy effectiveness of computing systems, see [Annex A](#).

6.2 Overview of holistic framework

This subclause provides an overview of holistic framework. When evaluating energy effectiveness of computing models, it is generally necessary to consider multiple energy effectiveness metrics due to the heterogeneity of components comprising the computing models. For example, the energy effectiveness of computing node may be expressed in terms of CPU speed over consumed energy, whereas the effectiveness of network node may be expressed in terms of network bandwidth over consumed energy. Holistic energy effectiveness evaluation means that the evaluation is performed considering multiple energy effectiveness indices simultaneously. The effectiveness indices may be homogeneous or heterogeneous among them. The holistic methods allow easy comparison of holistic energy effectiveness among multiple computing models or temporal trend. Also, it is possible to understand how the change of values in energy effectiveness metrics contributes to the overall energy effectiveness. For example, if the operator of a data centre wants to increase entire data centre energy effectiveness by double, it is possible to calculate how much improvement should be performed for each metric. During the holistic evaluation, it is possible to consider the characteristics of each computing model when estimating the holistic energy effectiveness using multiple energy effectiveness metrics. It is noted that depending on the service level objectives and agreements between customers and service providers, the performance and availability of services may vary. These service level objectives may affect the energy consumption of customer services. Thus, it is desirable to take account of customer expectations in order to provide more accurate energy effectiveness estimation.

6.3 Considerations for evaluating the applicability

This subclause provides considerations for developing and evaluating the applicability of energy efficient technologies for computing systems during sustainable ICT product during use stage of life cycle. In the computing models, energy is consumed during computing, data transport or data storing operations. In other words, not only data transport, but also data manipulation causes energy consumption. Moreover, the rapid deployment of mobile devices enables users to use multiple personal devices and maintain synchronized data among the devices. Therefore, user's activities can cause more effects to the energy consumption of ICT infrastructure than legacy environments. Current computing models impact energy requirements in multiple systems. Also, use of enterprise as computing, data source and maintenance requires data management across multiple systems. Low power, low capability end devices require increases in enterprise services and availability on resilient networks, servers and storage, and low capability end devices increases data centre and network energy requirements.

Therefore, the evaluation of computing models needs to consider the following issues:

- consideration for application's impact on ICT infrastructure in terms of data transport, computation and data storing;
- consideration for application's characteristics when executing the application. For example, users may execute application in a local node that has relatively, but does not require data transport over network or may execute the application at remote cloud node that has high performance, but needs data transmission over networks.

6.4 Examples of energy effectiveness evaluation

As investigated in [Clause 5](#), the energy effectiveness of end-to-end communication is represented as a ratio of total amount of data transferred, processed and stored to total energy consumption. In this subclause, it is described how to evaluate end-to-end energy effectiveness by using examples. [Figure 3](#) shows an example scenario for end-to-end energy effectiveness evaluation. In this example, a client device sends data traffic to the server in a data centre, and the sever stores the received data to an external storage systems.




Figure 3 — Example scenario for end-to-end energy effectiveness evaluation

This document considers the following assumptions.

B = is the amount of data traffic (Mbytes) that client device transmits to the server in the data centre.

T = is the data traffic transmission time.

S = is the data traffic storing and retrieving time of storage.

ΔS = is the time interval between storing and retrieving the data from storage.

N = is the number of hop counts in transport network.

P_{AVG_client} = is the average power consumption rate for client device during T .

$P_{AVG_router,j}$ = is the average power consumption rate for router j during T .

P_{AVG_server} = is the average power consumption rate for server during T .

$P_{AVG_stor_acs}$ = is the average power consumption rate for storage to store or retrieve data B during S

$P_{AVG_stor_maintain}$ = is the average power consumption rate for maintaining data in storage for ΔS .

It is noted that in case that P_{AVG_client} , $P_{AVG_router,j}$, P_{AVG_server} , $P_{AVG_stor_acs}$ and $P_{AVG_stor_maintain}$ are not available, P_{MAX_client} , $P_{MAX_router,j}$, P_{MAX_server} , $P_{MAX_stor_acs}$ and $P_{MAX_stor_maintain}$ may be used as an approximation.

P_{MAX_usr} = is the maximum power consumption rate for user equipment.

$P_{MAX_router,j}$ = is the maximum power consumption rate for router j .

P_{MAX_server} = is the maximum power consumption rate for server.

$P_{MAX_stor_acs}$ = is the maximum power consumption rate for accessing storage.

$P_{MAX_stor_maintain}$ = is the maximum power consumption rate for maintaining data in storage.

According to the description of end-to-end energy effectiveness presented in [Clause 5](#), the maximum energy effectiveness, η_{E2E} , can be expressed as follows. It is noted that this document refers only to computing model dedicated to a single data flow. Also, server is assumed as a dedicated server to a single application.

$$\begin{aligned}
\eta_{E2E} &= \frac{B_{E2E}}{E_{E2E}} \\
&= \frac{B}{E_{\text{client_device}} + \sum_{j=1}^N E_{\text{router}_j} + E_{\text{server}} + E_{\text{storage}}} \\
&= \frac{B}{\sum_0^T P_{\text{AVG_client}} + \sum_{j=1}^3 \sum_0^T P_{\text{AVG_router}_j} + \sum_0^T P_{\text{AVG_server}} + \sum_0^S P_{\text{AVG_stor_acs}} + \sum_0^{\Delta S} P_{\text{AVG_stor_maintain}}}
\end{aligned}$$

Figure 4 shows another example for e-mail exchange among users.

Similar to the previous example, the following assumptions are considered.

B = is the size of an e-mail message in Mbytes including attachments that user 1 sends to user 2 and user 3. It is assumed that the users use an e-mail server located at the data centre.

T = is the data traffic transmission time.

S = is the data traffic storing or retrieving time of storage.

ΔS = is the time interval between storing and retrieving the data from storage.

N_{usr1} = is the number of hop counts in transport network between the user 1 and the mail server in data centre.

N_{usr2} = is the number of hop counts in transport network between the user 2 and the mail server in data centre.

N_{usr3} = is the number of hop counts in transport network between the user 3 and the mail server in data centre.

$P_{\text{AVG_client}}$ = is the average power consumption rate for user equipment during T .

$P_{\text{AVG_router}_j}$ = is the average power consumption rate for router j during T .

$P_{\text{AVG_mail_svr}}$ = is the average power consumption rate for mail server during T .

$P_{\text{AVG_stor_acs}}$ = is the average power consumption rate for storage to store or retrieve data B during S .

$P_{\text{AVG_stor_maintain}}$ = is the average power consumption rate for maintaining data in storage for ΔS .




Figure 4 — Example scenario for e-mail exchange among users

According to [Figure 4](#), it is possible to identify three data transport paths in this e-mail exchange example.

- ① Deliver an e-mail message to mail server and store the e-mail to storage that are located at data centre.
- ② Deliver the e-mail message from the mail server to user 2 equipment using routing path 2.
- ③ Deliver the e-mail message from the mail server to user 3 equipment using routing path 3.

Therefore, by summing up the energy consumption of the three data transport paths, it is possible to approximately calculate the total energy consumption of the e-mail exchange scenario.

Similar to the previous example, the energy effectiveness of the example scenario can be calculated as follows.

$$\begin{aligned}
\eta_{E2E} &= \frac{B_{E2E}}{E_{E2E}} \\
&= \frac{B}{\sum_i E_{\text{client_device},i} + \sum_j E_{\text{router},j} + E_{\text{server}} + E_{\text{storage}}} \\
&= \frac{B}{\sum_{i=1}^3 \text{Energy Consumption of each routing path } (i) + \sum_{i=1}^1 \text{Energy Consumption of storage equipment}} \\
&= \frac{B}{\sum_0^T P_{\text{AVG_usr1}} + \sum_{j=1}^3 \sum_0^T P_{\text{AVG_router},j} + \sum_0^T P_{\text{AVG_mail_svr}} + \sum_0^S P_{\text{AVG_stor_acs}}} \\
&= \frac{B}{\sum_0^T P_{\text{AVG_client}} + \sum_{j=1}^3 \sum_0^T P_{\text{AVG_router},j} + \sum_0^T P_{\text{AVG_server}} + \sum_0^S P_{\text{AVG_stor_acs}} + \sum_0^{\Delta S} P_{\text{AVG_stor_maintain}}} \\
&= \frac{B}{\sum_0^T P_{\text{AVG_client}} + \sum_{j=1}^3 \sum_0^T P_{\text{AVG_router},j} + \sum_0^T P_{\text{AVG_server}} + \sum_0^S P_{\text{AVG_stor_acs}} + \sum_0^{\Delta S} P_{\text{AVG_stor_maintain}}}
\end{aligned}$$

7 Guidelines for determining the energy effectiveness of a computing model

This clause provides the basic steps to evaluate energy effectiveness of a computing model in terms of requirements, procedure, metrics and measurements. The general procedure for energy effectiveness determination is as follows.

- Identify the target computing model.
- Identify the components of the target computing model. The components can be explained in terms of computation, data transport and data storage.
- Determine the energy effectiveness metrics for the components in the target computing model.
- Determine the energy effectiveness of the target computing model from end-to-end perspective.
- Determine the energy effectiveness of the target computing model using the holistic method described in [Clause 6](#).

Annex A (informative)

State of affairs for improving energy effectiveness of computing systems

A.1 Taxonomy

Energy-effectiveness certainly is an important issue in networks and computing systems. This annex investigates various potential technologies for improving energy effectiveness of computing systems so that it provides a taxonomy of energy consumption points for computing systems in the computing models investigated in [Clause 5](#).

As the necessity for considering energy consumption in computing and network systems increases, a lot of research has begun to investigate energy awareness in networks and systems. This annex presents taxonomy of energy consumption points for computing and network systems based on the computing models investigated in [Clause 7](#).

A.1.1 Classification criteria

It is generally known that the energy consumption point of computing models can be divided into computing systems and networks. Thus, separate criteria for computing systems and networks are described.

This document refers the criteria for classifying energy effectiveness of computing systems and networks presented in Reference [2]. The general criteria for classifying energy effectiveness of networks include timescale, scope of required information, networking layer, input process and approach.

A first important criterion deals with the timescale of the decisions involved by the green strategy. As described in Reference [2], timescales on the order of nanoseconds to microseconds can be applied to CPU and instruction level, which is relevant in the computer and software architecture levels so these timescales concern individual components of a single system. Timescales on the order of micro to milliseconds deal with the system layer. At these timescales, actions may be taken between consecutive packets of the same flow, possibly involving several components at the same time, but likely confined within a single system. Larger timescales, on the order of one second and above, allow instead the action to span between multiple systems, possibly involving coordination of such systems as well. In order to present simple classification, it is possible to divide timescales into online and offline. Online requires actions to be taken during the operation of systems, which takes typically less than several seconds, whereas offline requires more time and the actions to be taken before runtime of systems as, for instance, during the system design process.

The second criterion concerns the type and amount of components or systems involved. In other words, it can be divided into local or global depending on the scope of the information required to take a decision. Local strategies will require information that pertains to a single system, or single network link, whereas global strategies will require information that pertains to a set of systems and links.

The third criterion is related with networking layer. Since most network equipment is implemented with layering concept, it is possible to classify energy consumption points according to the networking layer to which they apply. Based on the TCP/IP protocols stack, each energy consumption point can either be classified into data link, network, transport and application layers, or may require cross-layer interaction.

Another classification criterion comes from the analysis of the input process that drives the decision taken in the solutions. The decision may be taken based on the instantaneous situation, on the historical pattern or on the forecast. In the case of online solutions, all the three kinds of decision are possible. For the offline solutions, historical pattern and forecast based analysis are only applicable.

Finally, another criterion related to the evaluation methods of the proposed technology is required. Typical examples of the evaluation methods are discrete event simulation, hardware/software prototyping or formal models with analytical or numerical solution. All methods have their respective advantages and disadvantages. Theoretical analysis and simulation can be performed with low cost and short time compared to prototype, hardware or real deployment. However, the latter approaches may provide more precise evaluation and a higher level of maturity of a specific research field. [Table A.1](#) shows the summary of the classification criteria^[2].

Table A.1 — Summary of the classification criteria for energy efficient technology

Criterion	Category	Meaning
Timescale	Online, offline	Defines the update frequency of the policy.
Scope	Local, global	Influences the volume of communication required to reach the objective.
Networking layer	Link, network, transport, application, cross-layer	Individuate which entities shall collaborate.
Input	Instantaneous, historical, forecast	Defines learning and adaptation capabilities of the algorithm.
Evaluation approach	Traffic analysis, theoretical modelling, simulation, hardware and software prototyping	Flavour of the study, also correlates with the level of maturity of the work.

A.1.2 Taxonomy of energy efficient technologies

This subclause presents taxonomy of energy efficient technologies for computing models according to the classification criteria investigated in [A.1.1](#). As discussed, computing models can be divided into computing systems and networks. Thus, the taxonomy consists of three categories, namely computing systems, networks and storages, respectively.

Technologies for computing systems are categorized as follows.

— Hardware-level technologies

- Dynamic voltage scaling: a technology that allows the voltage supplied to some computer components to be raised or lowered dynamically according to power, performance and thermal requirements.
- Dynamic frequency scaling: a technology that allows the CPU clock speed to be automatically adjusted according to power and thermal requirements.
- Link power management: technologies that allow parts of connected devices in a computer or connected to a computer to enter low power states when idle.
- Energy efficient display: a computer display or monitor that uses energy saving technologies such as organic light emitting diode (OLED) or techniques to reduce power consumption when in sleep mode.
- Energy efficient power supply: a power supply which meets specified energy targets set by organizations such as 80 PLUS.

- Energy efficient battery charger: a battery charger for cellular phones, laptop PCs or tablets that incorporates energy saving techniques.
- Software-level technologies
 - Workload consolidation: a technique that allows running multiple tasks on the same physical machine in order to reduce the number of nodes that are switched on. A key component of systems that aim to consolidate workloads is to monitor and estimate the load posed by user applications or estimate the arrival of user requests.
 - Energy-aware task scheduling: energy aware task schedulers have three types, namely offline scheduling based on a prior task information, online scheduling which is purely dynamic and hybrid approaches, including an offline phase where the slack is greedily absorbed and dynamic algorithms operating in the online phase.
 - Virtualization: creates a virtual machine that acts like a real computer with an operating system and improves energy efficiency.
 - Virtual machine migration: moves a running virtual machine or application between different physical machines without disconnecting the client or application.

Technologies for networking systems are categorized as follows.

- Adaptive layer-2 technology: most current types of network equipment show a constant power consumption profile largely irrespective of their actual utilization. The computing world on the other hand has long embraced methods to approximate energy-proportional computing. In the networking world, there are initial steps into energy-proportional communications. Adaptive link rate technology is such a step where the energy use changes with link utilization.
- Network interface proxy technology: proxying describes technologies that maintain network connectivity for other devices so that these can go into low power sleep modes. This mainly targets the reduction of unnecessary energy waste through edge devices.
- Energy-aware infrastructure technology: in order to achieve a further reduction in energy use, coordination and management of larger parts of the network appears to be a promising idea. Energy-aware infrastructure describes a class of techniques to this end. Energy-aware routing is one example that falls into this category. Energy-aware routing makes use of the fact that traffic follows certain patterns. Based on this knowledge, in times where network traffic is low, a number of routers can be put to sleep while the network as a whole still preserves connectivity and an adequate service level. Requirements for an energy-aware control planes are outlined in Reference [5].

Technologies for storage systems are categorized as follows.

- Node-level techniques: this category can be further classified into caching, prefetching, I/O scheduling and disk spin up/down. The purpose of caching and prefetching techniques is to store frequently used data in the fastest and more energy efficient memories of the storage hierarchy, thus decreasing requests to main storage. Similarly, disk scheduling techniques reschedule I/O operations in order to prolong disk standby times. The longer standby times produced by these techniques increase the effectiveness of disk power-saving features, which are exploited by disk spin up/down solutions.
- Distributed techniques: this category includes data placement and consolidation. Data placement is layout strategies, which help achieve resource consolidation. The purpose of consolidation is to maximize power proportionality by having a smaller number of nodes working at a higher utilization rate, instead of having many at a lower utilization rate, which is highly inefficient.

A.2 Practice cases for energy effectiveness improvement

This subclause provides survey of current known technologies for developing energy efficient computing systems.

A.2.1 Computing system-level technology

There exist many technologies for improving the energy effectiveness of computing systems in terms of hardware and software. Some of the examples are as follows.

- Hardware-level technologies: the technologies in this category are mainly targeted to improve the energy effectiveness of hardware components of a single computing system. For example, chip fabrication, CPU design, clock and circuit at chip, memory technology and so on.
- Software-level technologies: the technologies in this category improve energy effectiveness of single system. For example, idle/sleep model management of a system, efficient cooling system control, contents caching, traffic shaping, etc.

A.2.2 Networking system-level technology

A.2.2.1 Adaptive layer-2 technology

There exist many research proposals to adapt the link rate to network traffic pattern by either turning off links during idle periods which is usually referred to as sleeping mode, or by reducing the line rate during low utilization period which is known as rate switching. In both cases, the link rate is adapted to match the actual network usage, hence achieving a link utilization proportional energy consumption rate. This technologies aim at bringing the energy consumption of a link from the initial constant worst case down to a curve closer to the ideal proportional case.

An early work regarding adaptive link rate based energy efficient systems only considers two states of operation: an idle mode and a fully working mode[6][7][8]. The difficulty in this case consists in finding the desired compromise between system reactivity and energy savings. A recent survey covering sleeping mode from an algorithmic point of view can be found in Reference [9]. In the pioneering work[6], the authors let the nodes decide on their interfaces status by measuring packet inter-arrival times and considering if this interval is long enough to justify an effective energy saving between two consecutive frames. As the effectiveness of such a method is directly tied to the inter-arrival distributions, the authors analysed a traffic trace to determine whether such an approach would be effective or not in practice[2].

The early research work introduces a number of questions, addressed by subsequent research. First of all, different types of sleeping mode are possible for an interface, depending on the technology. The interface may be in a deep idle state and drop the packets arriving during the sleeping period[8], be fully awakened by every packet reception, use a buffer to store the packets received during the sleeping intervals processing them when it wakes up or use a shadow port that may handle the packets on behalf of a cluster of sleeping ports[10]. However, even semi-sleep state does have a price in terms of energy. First, any idle state in which packets can be detected requires some electronics to be active and thus, consumes a small amount of energy, as well as powering any shadow port does. Hence, waking up the sleeping interface at every packet arrival reduces the packet delay and loss, but it will also reduce the energy savings[2].

In Reference [7], the authors proposed a two-state model of the sleeping mode strategy. The first state corresponds to the regular operation mode, and the second one to the energy saving mode. The first transition, from the energy saving mode to the operational mode, takes time and generates an energy consumption spike. The second transition, from the regular working mode to the energy saving mode, is supposed to be instantaneous and spike-less. This approach results in a simple model that can easily be extended to include more than two states, modelling, for instance, rate switch strategies.

Besides the choice between an idle and a working mode, most of current research proposes a wider range of possibilities through the use of several transmission rates[11][12][13] to which different energy consumption figures correspond. For example, Ethernet defines several transmission rates, from 10 Mbps to 10 Gbps and, etc. The authors of Reference [11] show that there is a non-negligible difference in the interface energy consumption by configuring different data rates. For example, an increase of the data rate of a PC end system network interface card (NIC) from 10 Mbps to 1 Gbps results in an increased energy consumption of about 3 W, which represents about 5 % of the overall system

energy consumption. For regular switches, the same throughput shift results in a per interface energy consumption increase of about 1,5 W per link. The same authors of Reference [11] proposed successive refinement of the rate control policies based only on the current system state [12], or on an historical analysis [13]. From a high-level point of view, selecting the proper data rate among a limited set of possibilities can be translated in an integer linear program whose objective is to minimize the overall energy consumption given that all the data inputted at the network is forwarded, which is known to be NP-hard [2].

Regarding the standardization activities, IEEE energy efficient Ethernet is a step into this direction. The respective task force has devoted considerable efforts in this area and has published IEEE 802.3az-2010 (amendment to IEEE 802.3-2008) [3]. The amendment describes modifications to existing physical layer specifications in order to make Ethernet more energy efficient. The amendment covers various technologies such as 10BASE-T or 100BASE-TX. It also proposes a new low power idle (LPI) mode and related mechanisms and protocols in order to energy efficiently manage Ethernet links.

A.2.2.2 Network interface proxy technology

Network proxy refers to a set of mechanisms dedicated to put network interfaces and networked nodes into energy saving idle mode. Energy consumption in idle mode is less than active mode in general, so the longer the idle periods is, the higher the achievable energy saving can be. Network proxy is a technology that delegates some of networking functions in networked nodes to other network elements such as external nodes or other modules in the nodes. The network proxy enables network nodes to maintain network connectivity during idle mode. When a host wants to enter energy saving idle mode, the host delivers its network status and state to a network proxy and goes into idle mode. Then, the network proxy responds to periodic messages on behalf of the host in idle mode. If the proxy receives a message that it cannot process, it sends a wake-up message to the host so that the host can process the message after wake-up.

According to the survey, even though users are idle, background network traffic is nevertheless received and needs processing preventing thus PCs from going in sleeping mode. Also, it is known that most of the incoming traffic received by a PC network interface during otherwise idle periods can simply be dropped or does not require more than a minimal computation and response. For instance, most broadcast frames or traffic related to port scanning may simply be ignored. Usual exchanges, such as address resolution protocol (ARP) processing, internet control message protocol (ICMP) echo answering or dynamic host configuration protocol (DHCP) rebinding, are simple tasks that could be easily performed directly by the network interface. The idea behind network interface proxying consists in delegating the processing of such traffic. Processing can imply plain filtering or may require simple responses (e.g. in the case of ARP, ICMP, DHCP) or even more complex task. Such tasks can be delegated from the CPUs of end devices to a number of different elements either local elements such as low-energy processor on-board of the network interface card (NIC) of the same device or to an external element in a LAN environment [2].

NIC proxy implements light processing function of the received packets in the NIC. The NIC may drop the periodic protocol message exchange and handle the traffic requiring minimal computation, while the full system will be woken up only when non-trivial packets needing further processing are received. The NIC proxy allows energy saving through power down the end hosts without losing their network connectivity. According to Reference [2], the NIC proxy technique may be applied to more than 90 % of the received packets on a host during idle periods.

External proxy is offloading traffic processing to an external system within a LAN, so that the proxy acts for a number of end-devices. It can feature a more efficient CPU and thus, offload the end-host from an even higher number of network maintenance tasks. Delegation of ARP processing is a typical example of external proxy. For example, a switch acts as a proxy for ARP traffic, allowing the target machine to sleep at least until data traffic is sent. Energy-aware proxies are instead implemented in Reference [10] as a modular routers. The authors implemented four different kinds of proxies, of increasing complexity, showing that although the potential energy saving is considerable, nevertheless trivial approaches are not sufficient to fully exploit the potential saving. Indeed, while broadcast traffic is easily filtered, a significant implementation effort is needed to properly handle unicast traffic. Finally, it is noted that all

the above work do not take into account the residential environment, where set-top-boxes are likely to offer opportunities for external proxy functionality[2][14].

Regarding the standardization activities for network proxy, ECMA has published a proxying document. It specifies maintenance of network connectivity and presence by proxies to extend the sleep duration of hosts. It specifies capabilities that a proxy may expose to a host, information that must be exchanged between a host and a proxy, and required and optional behaviour of a proxy while it is operating, including responding to packets, generating packets, ignoring packets, and waking the host. However, this document does not specify communication mechanisms between hosts and proxies and extension or modification of the referenced specifications and support of security and communication protocols such as IPsec, MACSec, SSL, TLS, Mobile IP, etc.[4]. The ECMA specification has been adopted by ISO/IEC JTC1 and has also been published as ISO/IEC 16317:2011[15].

A.2.2.3 Energy-aware infrastructure

The mechanisms discussed in the previous subclauses involve local decisions, through a single device or a very small set of collaborative devices. While these techniques alone offer non negligible energy saving, further improvement can be expected from a reasonable amount of collaboration between individual devices, sharing a wider knowledge on the system state. One approach is to schedule resource on/off periods according to the received request and evaluate the energy savings by applying a resource energy management method on a historical request record[16]. Another incremental approach is proposed in Reference [17], which considers the automatic adaptation of the link rate from the global viewpoint of a whole backbone network. In detail, traffic is reshaped into bursts at the network edge by arranging all packets destined to the same egress router to be contiguous within the bursts. This approach adds end-to-end delay, but only at the network ingress. Consequently, periods of activity and sleep alternate inside the network allowing fewer state transitions with respect to a simple opportunistic “wake on arrival” strategy. The work studies the incidence on the time spent in sleeping mode for different parameters: average network utilization, burst size and state transition time. The authors argue that this strategy does not add significant complexity to the network, although they do not propose strategies to determine when and for how long nodes should ideally sleep (notice that global coordination may be very hard to achieve). Moreover, the effects of the traffic shaping on the jitter are not analysed, even if they are expected to be important.

Other works instead advocate the use of clean-slate approaches, with a higher use of optical networks (e.g. Dense Wavelength Division Multiplexing). It is now admitted that optical switching is much more energy efficient, while offering an extremely large capacity. At the same time, these technologies still suffer from a lack of flexibility with respect to the electronics domain. A future challenge is probably to find efficient architectures combining both optical transport and packet processing, when needed.

Finally, the problem of introducing energy-awareness into the network design process is studied in References [18] and [19] from an operational research point of view. In more detail, Reference [18] introduces the energy consumption cost into the multi-commodity formulation of the design problem, together with the performance and robustness constraints. A similar approach is adopted in Reference [19], which evaluates also the trade-off between energy consumption and network performance, highlighting the fault tolerance characteristics of the different possible working points.

If much improvement is expected from the link layer, through link adaptations and proxying techniques, the network layer may also be involved in the reduction of the energy expenditure. Following the resource-consolidation principle, energy-aware routing generally aims at aggregating traffic flows over a subset of the network devices and links allowing other links and interconnection devices to be switched off. These solutions shall preserve connectivity and QoS, for instance, by limiting the maximum utilization over any link or ensuring a minimum level of path diversity. Flows aggregation may be achieved, for example, through a proper configuration of the routing weights. Formally, energy-aware routing is a particular instance of the general capacitated multi-commodity flow problems[20], and thus, falls into the set of the offline solutions concerning dimensioning.

Energy-aware routing has been first evoked in the position paper[21], but just as a hypothetical working direction (under the name of “coordinated sleeping”) by taking the example of two parallel routers acting on the boundaries of an AS. The possibility of coordinating the sleeping periods of these two

routers is discussed, speculating on the impacts on fault tolerance and the required changes to routing protocols. On the one hand, OSPF considers, by default, sleeping links as faulty and requires an update of the shortest paths, which requires time consuming computation. On the other hand, IBGP may suffer from routes oscillation and see occasionally forwarding loops. The paper indicates a possible solution to this problem through the use of different pre-computed solutions or in the presence of a unique centralized omniscient decision point.

A.2.3 Storage system-level technology

There exist many technologies for improving the energy effectiveness of computing systems in terms of hardware and software. Some of the examples are as follows.

- Virtualization, tiering and application alignment: virtualization of storage enabled multiple systems to share and access a single storage device. The storage infrastructure included both SAN and NAS environments. Virtualized storage environments enable tiered storage and transfers data amongst virtual storage machines with relative ease. Tiering and application alignment had impact on performance by increasing utilization and scalability, enabling multi-vendor sourcing and simplifying management. Tiering strategies reduced overall ownership cost of the storage. Storage-medium allocation improvement was done through NAS/SAN virtualization, system to application mapping and alignment and data migration.
- Capacity management: capacity management had improved the utilization rate of storage devices by implementation of various techniques such as thin provisioning, fabric unification, storage reclamation and capacity management reports and metrics.
- Data management: data management technologies reduced the volume of data to be stored and restricted the capacity growth through various techniques like writable snapshots, deduplication, next-generation backup and recovery, etc.
- Solid state disk (SSD): persistent storage devices that use semi-conductor memory technologies. SSD have no moving parts and use less power per unit of performance than most rotating media hard disks.
- Data reduction: data reduction technologies such as compression, deduplication and thin provisioning reduce storage consumption. Deduplication reduces redundancies that consume disk space unnecessarily. Thin provisioning enables more efficient disk space utilization by not reserving capacity from the storage allocation pool before there is actual data to be stored.
- Dynamically adjusting disk rotation speed (DRPM): it allows dynamic modulation of the disk spin speed so DRPM adjusts disk rotation speed to match the required performance.
- Energy efficient file systems: file systems are used for organizing data in data storage devices. Distributed file systems are widely employed in large scale computing systems. In the context of high performance computing, parallel file systems are widely used. While most of the file systems have been designed in order to maximize performance, there exists a few of them having been designed in order to target the problem of energy efficiency^{[42][43]}.