low Energy COnsumption NETworks

DELIVERABLE D6.5

BENCHMARKING AND PERFORMANCE EVALUATION RESULTS

Grant Agreement Number:	**258454**
Project Acronym:	**ECONET**
Project Title:	**low Energy COnsumption NETworks**
Funding Scheme:	Collaborative Project
Starting Date of the Project:	01/10/2010 *dd/mm/yyyy*
Duration:	39 months
Project Coordinator:	Name: Raffaele Bolla Phone: +39 010 353 2075 Fax: +39 010 353 2154 e-mail: raffaele.bolla@unige.it

Due Date of Delivery:	M39 *Mx* (31/12/2013 *dd/mm/yyyy*)
Actual Date of Delivery:	28/01/2014 *dd/mm/yyyy*
Due Date of Revision Delivery:	M44 *Mx* (31/05/2014 *dd/mm/yyyy*)
Actual Date of Revision Delivery:	31/05/2014 *dd/mm/yyyy*
Workpackage:	**WP6 – *Integration, experiments and performance evaluation***
Nature of the Deliverable:	R
Dissemination level:	PU
Editors:	MLX, CNIT, TELIT, ALU, TEI, LQDE, INFO, ETY, NVR, GRNET
Version:	2.0

List of the Authors

MLX	MELLANOX TECHNOLOGIES LTD - MLNX
Lavi Koch	
CNIT	CONSORZIO NAZIONALE INTERUNIVERSITARIO PER LE TELECOMUNICAZIONI
Roberto Bruschi, Alessandro Carrega, Franco Davoli, Paolo Lago, Chiara Lombardo, Sergio Mangialardi, Fabio Podda, Mirko Rubaldo, Matteo Repetto, Antonio Cianfrani, Marco Polverini, Luca Chiaraviglio, Nanfang Li	
TEI	ERICSSON TELECOMUNICAZIONI
Renato Grosso	
ALU	ALCATEL-LUCENT ITALIA S.P.A.
Giorgio Parladori	
TELIT	TELECOM ITALIA S.p.A
Diego Suino	
GRNET	GREEK RESEARCH AND TECHNOLOGY NETWORK S.A.
Constantinos Vassilakis, Anastasios Zafeiropoulos	
NVR	NETVISOR INFORMATIKAI ES KOMMUNIKACIOS ZARTKORUEN MUKODO RESZVENYTARSASAG
Arpad Bakay	
ETY	ETHERNITY NETWORKS LTD
David Levi	
INFO	INFOCOM S.R.L.
Maurizio Giribaldi	

Disclaimer

Copyright

Table of Contents

1 Executive Summary

This document aims at reporting, discussing and analysing the experimental results obtained in the final demonstration of the ECONET project, and at demonstrating that the ambitious goal of reducing the energy consumption of wire-line network infrastructures by 50%-80% has been successfully achieved.

The demonstrator was set up in November-December 2013 in the Telecom Italia site of Largo Borgaro in Turin, Italy, and it included 15 physical prototypes of green network devices, in addition to ad-hoc network emulators realized by the ECONET consortium for evaluating the behaviour of green control plane protocols and algorithms.

The structure of the demonstrator has been extensively described in the D6.4 report [1], whose reading is warmly suggested before proceeding with the current document to better understand the benchmarking scenarios, as well as the general validity and significance of the results obtained.

The demonstrator has been specifically designed for representing an entire wire-line network infrastructure of a Telecommunication Operator (Telco) or Internet service Provider (ISP), along with customer premises and datacentre networks. For this reason, the 15 prototypes have been selected to represent the most typical nodes of the aforementioned network segments, spanning VDSL home gateways and DSLAMs, network business access to metropolitan (metro), core, and transport networks, up to datacentre network devices.

All the prototypes have been equipped with the most promising green data-plane technologies and solutions proposed in the project WP3. Then, all these data-plane technologies have been exposed through the Green Abstraction Layer (GAL), the interface designed in ECONET WP4 that is going to be accepted as an official standard by ETSI, to "Local Optimization Policies" and to the ECONET WP5 control and management planes. On these planes, a number of new green protocols (or green extensions to existing protocols) have been used (*i*) to interface energy-aware mechanisms to optimize the trade-off between performance and energy consumption on the entire network, and (*ii*) to expose green metrics and to make them controllable by network operators by means of extended network management and monitoring frameworks. In other words, the demonstrator is the real essence of the ECONET project, and can be literally considered as the tangible proof of all its baseline research concepts, development and integration efforts.

Within this very complete picture of the ECONET vision, a number of tests have been performed by emulating typical daily traffic fluctuations, letting the green mechanisms react to the different conditions and traffic levels, and measuring/monitoring their performance in terms of energy- and traffic-aware metrics.

Notwithstanding the maturity levels of some prototypes and a number of simplifications needed to build such a large-scale demonstrator (e.g., the presence of a significant share of legacy components), the raw experimental results clearly outline that ECONET technologies are able to provide by now an impressive reduction of power consumption in all the considered network segments.

It is estimated that, if the same setup of the demonstrator would be applied to a medium scale operator, energy saving exceeding 39% over the Business-as-Usual scenario could be achieved also in the presence of a non-negligible share of legacy hardware.

Thus, in order to evaluate the real potential impact of such technologies and to demonstrate that the project has achieved its target, the current document also includes an extended analysis to map the demonstration results on real telecom operator infrastructures in the presence of realistic operating

scenarios. The obtained results suggest that the ECONET technologies are capable, in the short-term, of an impact of 51% in energy saving with respect to the Business-as-Usual scenario, and of an impact of 77.6% in the long-term.

2 Introduction

The final ECONET Project Demo took part during the week of 9th Nov. 2013, in the Telecom Italia Test Plant laboratory, located in Turin, Italy. During the three weeks prior to this date, the ECONET partners involved in the final demonstration had shipped their devices to the TELIT lab premises, and the entire demo network was set up.

During the last week before the final demo, all the partners arrived at the TELIT site, and worked in collaborating effort to bring up and to integrate all the systems, as predefined in the D6.4 report [1] in detail. As made explicit by the D6.4 report, which defines the overall physical and logical architecture of the demonstrator, all the most significant data, achievements, technologies and solutions coming from the R&D Work Packages (WPs) have been fully integrated in the various network clouds, and carefully evaluated in the final demonstration.

It is worth underlining that the reading of the D6.4 report is warmly suggested before proceeding with the current document to better understand the benchmarking scenarios, as well as the general validity and significance of the results obtained.

Given the complexity of the testing scenario and the dimension of the demonstrator, these activities also included a large number of preliminary tests and checkpoints, aimed at validating single network clouds and the interoperability of the different systems, which were gradually added/connected to the demonstrator. The traffic generation, the daily profiles, and the generation time acceleration (i.e., the duration of time slots) were carefully tuned while enabling the green low power features, protocols and algorithms integrated in the prototypes.

This document aims at reporting, discussing and analysing the experimental results obtained in the final demonstration of the ECONET project, and at providing a tangible proof that the ambitious goal of reducing the energy consumption of wire-line network infrastructures by 50%-80% has been successfully achieved. Within this very complete and complex demonstrator composed of heterogeneous technologies working at various logical levels and network layers, a number of tests have been performed, by emulating typical daily traffic fluctuations, letting the green mechanisms react to the different conditions and traffic levels, and measuring/monitoring their performance in terms of energy- and traffic-aware metrics.

As a proof for the actual measurements taken, the demonstration activity is supported by a new section on the ECONET website (see Fig. 1), which reproduces the main results collected during the demonstration and provides access to a number of videos including live recordings of the actual demonstrations. These videos not only provide a zoom-in on the actual measurements observed during the demonstration, allowing a better understanding of single technologies, but they also constitute a means for disseminating the experience and the achievements of the ECONET project as widely as possible.

These live videos cover the following parts of the final demonstration:

- Optical Transport Cloud
- Metro Cloud
- Core Cloud
- Data Centre
- Network Connectivity Proxy
- Network management and monitoring

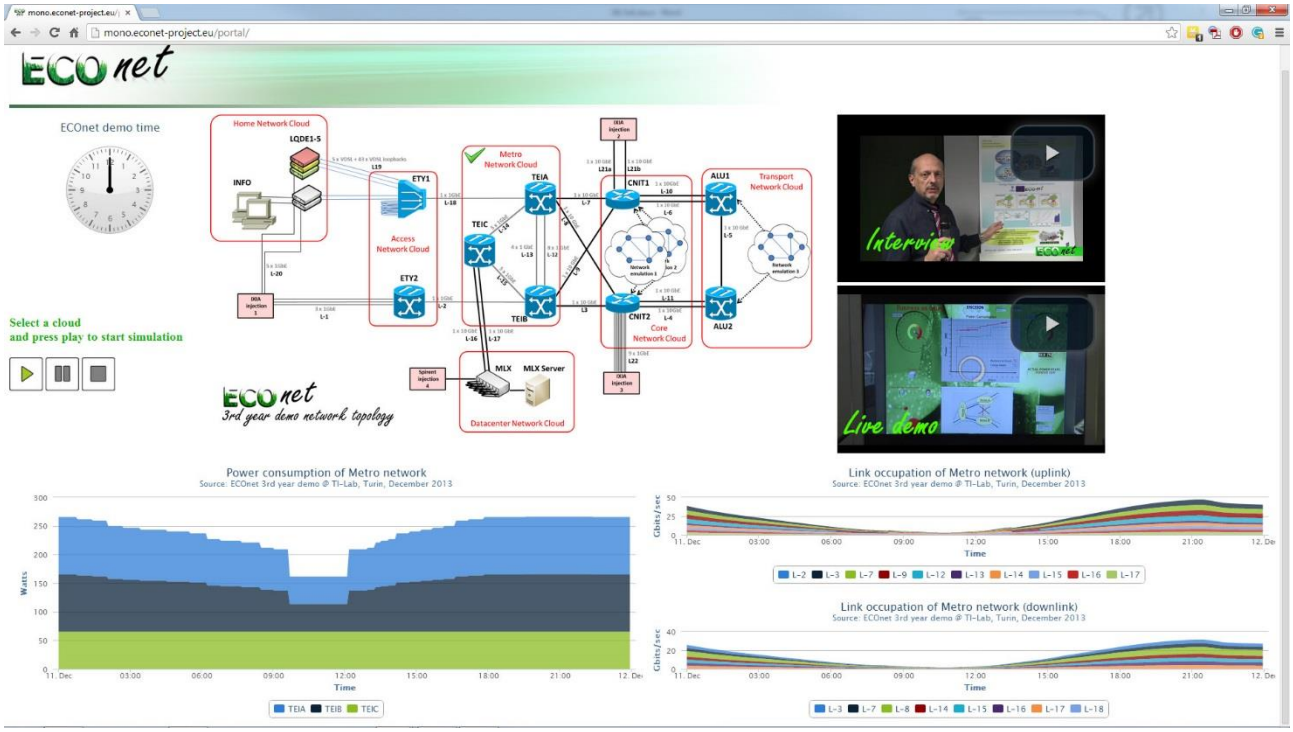

Fig. 1. The website section devoted to report the ECONET demonstration results and live videos.

1.1 Structure of the Document

In addition to an Executive Summary and the Introduction, the deliverable includes the following three sections:

Section 3, Experimental Results: this section briefly introduces the setup of the demonstrator and the main tests performed, and presents the experimental results obtained in the demonstration. In more detail, as first demonstration results are presented and discussed in an aggregated and high-level manner, then, detailed results for each single demonstration network cloud are reported and analysed. This section also includes the performance evaluation of the Network Connection Proxy when fully integrated with the demonstration network, describing the verification of the NCP behaviour, followed by its performance analysis, and concluding with its impact overhead.

Section 4, The expected impact of ECONET technologies: this section aims at re-elaborating the raw experimental data collected during the demonstration tests, and using them for estimating the real impact of ECONET technologies when applied to the entire infrastructures of network operators in realistic working environments. For this purpose, this section starts by describing the Energy Profile Model and the network infrastructures and scenarios used in the ECONET final demonstration. Then, by considering the results obtained in each demonstration cloud, it estimates expected power saving in each network segment.

Section 5, Conclusions: this section concludes the deliverable with a summary of the main results obtained in the final demonstration of the ECONET project.

3 Experimental Results

As defined in detail in the D6.4 report, all the experimentations have been carried out with the topology and the configuration depicted in Figs. 2 and 3, representing the paths of the traffic flows in day and night configurations, respectively.

Figs 4-11 show some pictures/screenshots taken during the experimentations, and show the physical setup of the demonstrator. In detail, Fig. 4 provides a global view of the prototypes composing, from left to right, the transport, core, datacentre and metro clouds, as they have been set up during the ECONET demonstration. The single racks hosting each cloud can be seen in Figs. 5-8.

Fig. 5 reports the transport network cloud, which reproduces the optical backbone of the network through two physical nodes and an emulation network. The stickers that are visible in this picture are useful to see how the ALU1 and ALU2 nodes are interconnected with links L-3, L-4, L-10, L-11 and L-5.

Fig. 6 contains the two instances of the DROP router that, in addition to the emulated virtual routers, compose the core network cloud, and the Ixia traffic generator. In this picture, the building blocks that form the CNIT1 and CNIT2 distributed routers are outlined as FEs (Forwarding Elements) and IEs (Interconnection Elements).

The bottom part of the rack shown in Fig. 7 hosts the servers running the virtual routers composing the emulated part of the core network, while the top host in the datacentre cloud is composed of a server equipped with a high-speed network adapter, and a high-speed interconnection switch.

The metro network cloud shown in Fig. 8 is composed by a typical ring topology of three metropolitan carrier Ethernet switches, which have been designed and provided by Ericsson.

Regarding the home and access clouds, the deployment exploited during the demonstration can be seen in Fig. 9. In detail, five home networks have been deployed each one composed by a VDSL HG and, optionally, some network hosts. The access cloud includes a VDSL DSLAM for residential or small business network terminations, and a Gigabit Ethernet switch for enterprise network access of wireless terminations back hauling.

To complete the setup description, it is worth mentioning the monitoring and OAM cloud, which was also physically present in the demonstrator. Fig. 10 shows a screenshot of the web GUI of the central energy-aware OAM system.

Finally, the synchronization of the physical and virtual building blocks with the day-night traffic profiles was guaranteed by a clock application, a screenshot of which can be seen in Fig. 11.

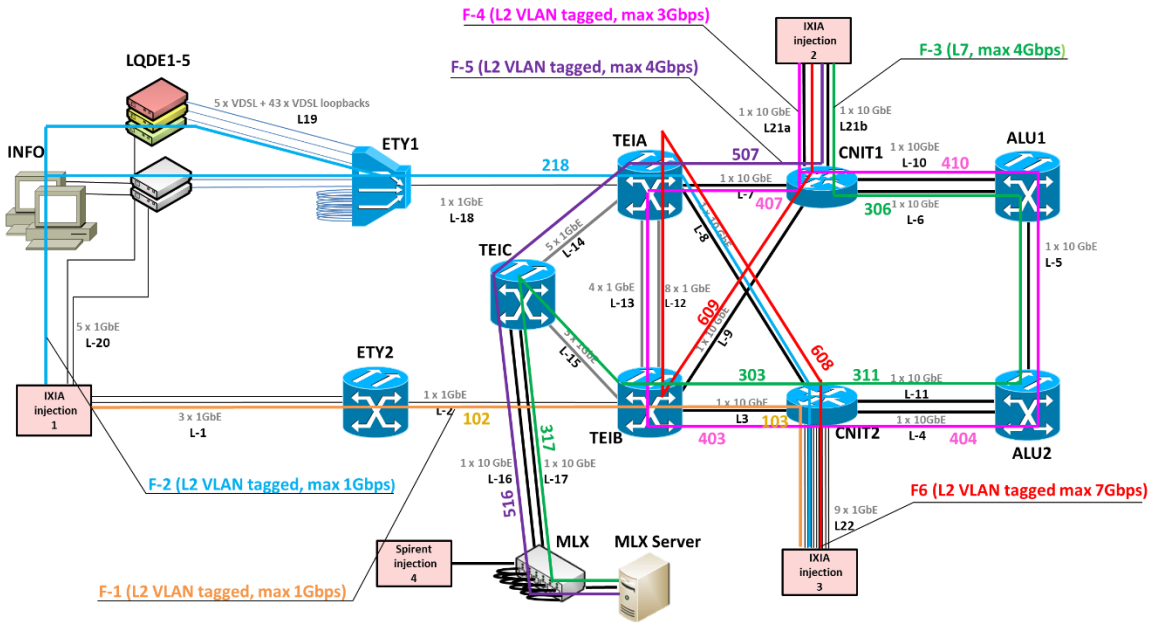

Fig. 2. Demonstration topology and traffic flows in the "Day Configuration".

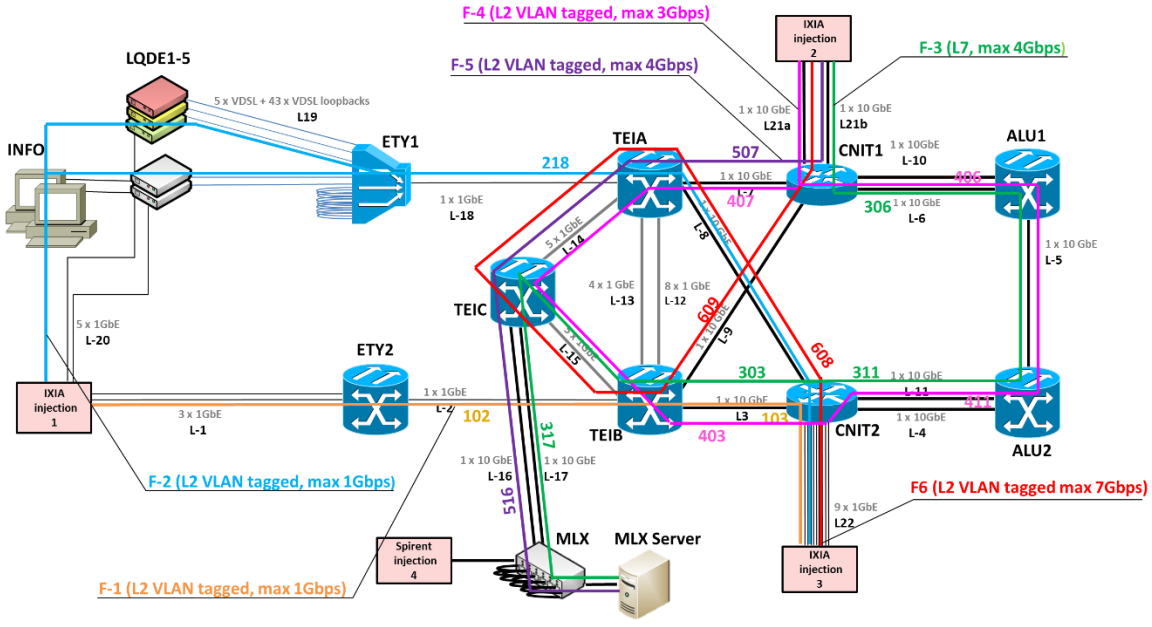

Fig. 3. Demonstration topology and traffic flows in the "Night Configuration".

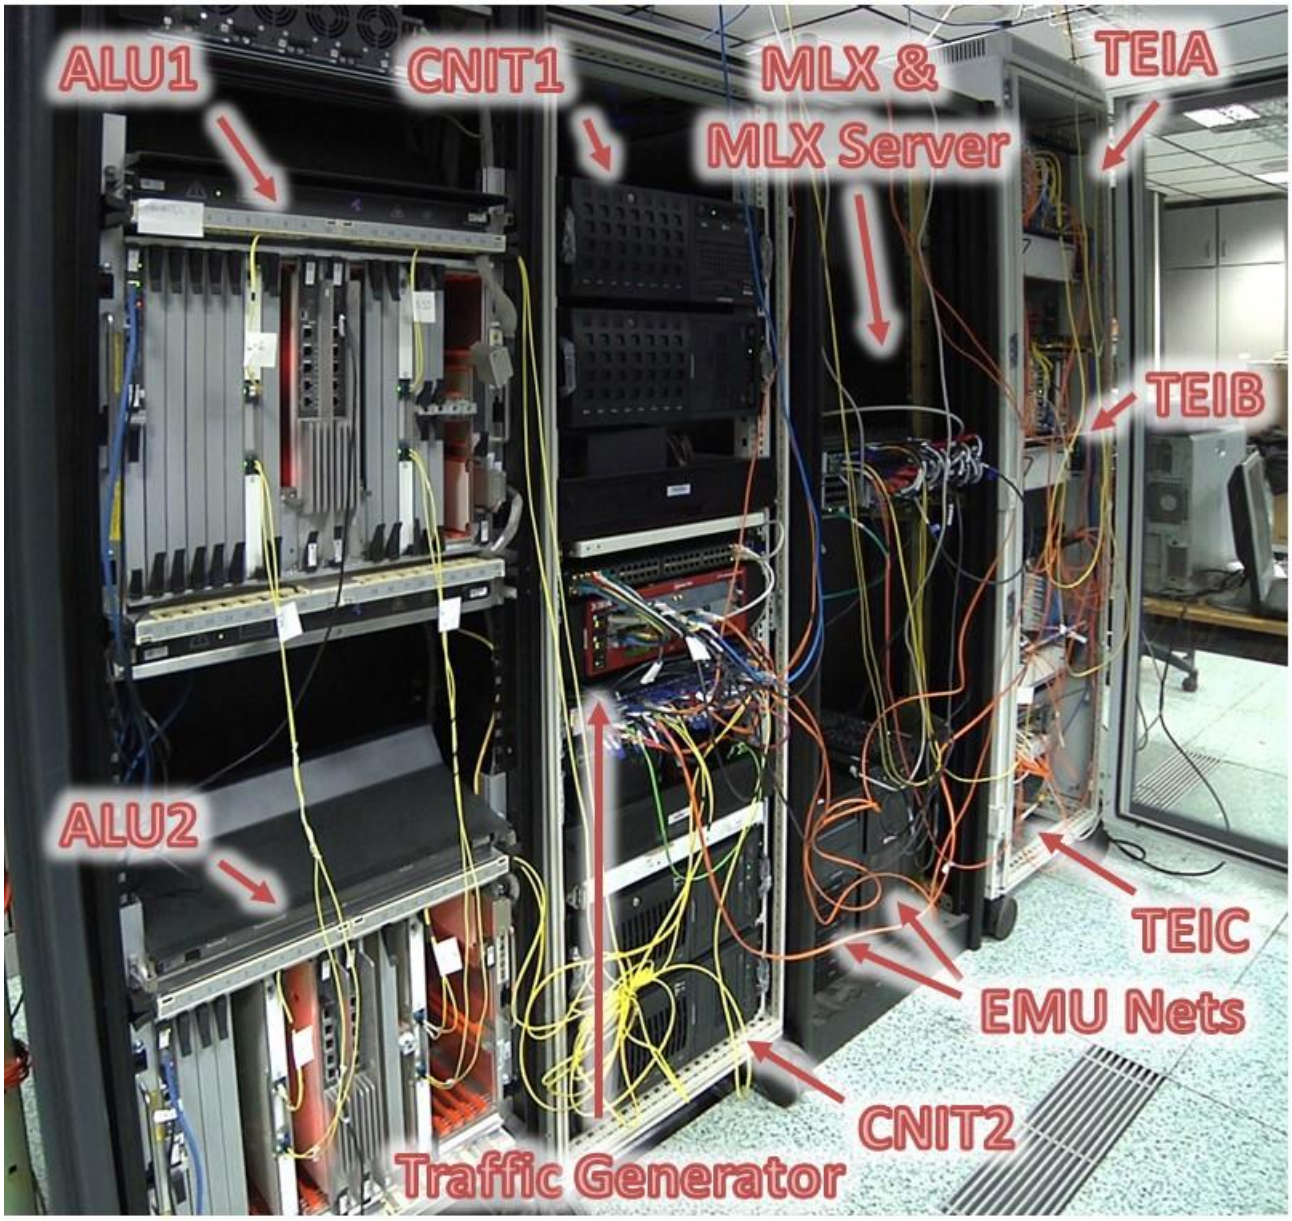

Fig. 4. Setup of the transport, core, metro and datacentre clouds during the ECONET demonstration.

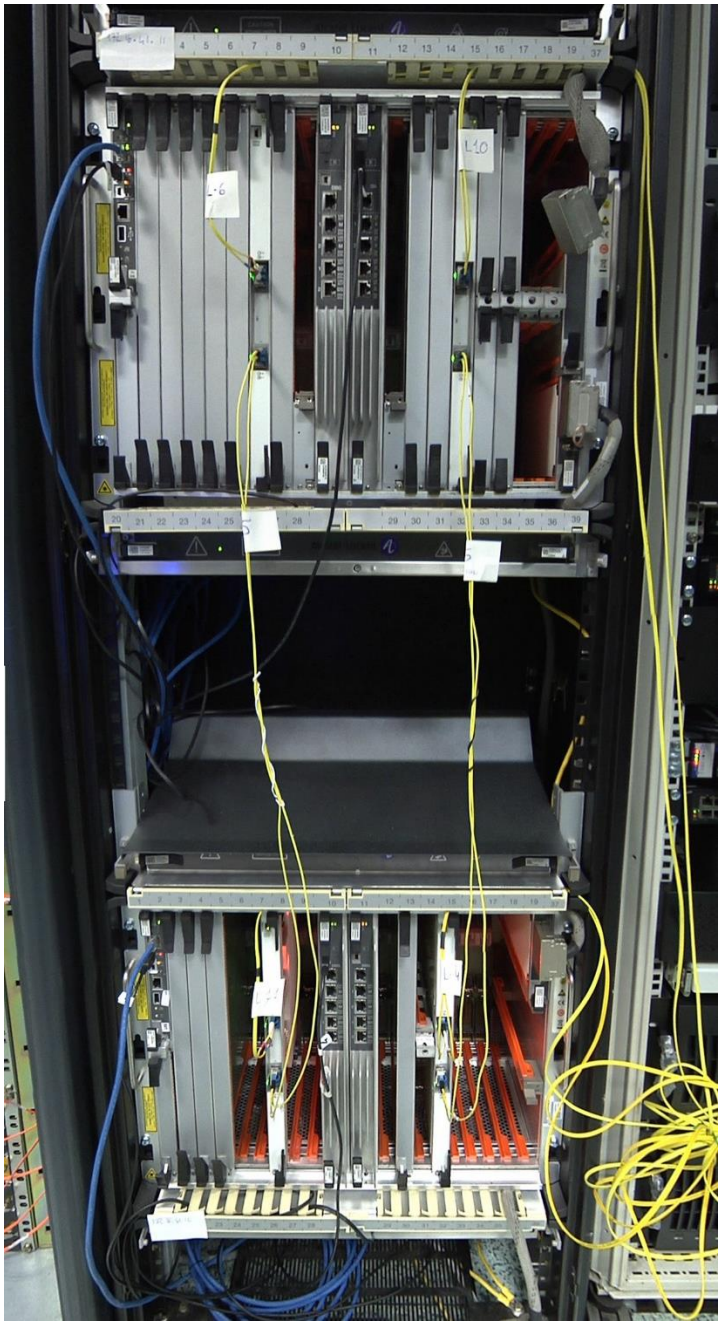

Fig. 5. The rack hosting the transport network with the ALU1 and ALU2 nodes interconnected with links L-3, L-4, L-10, L-11 and L-5.

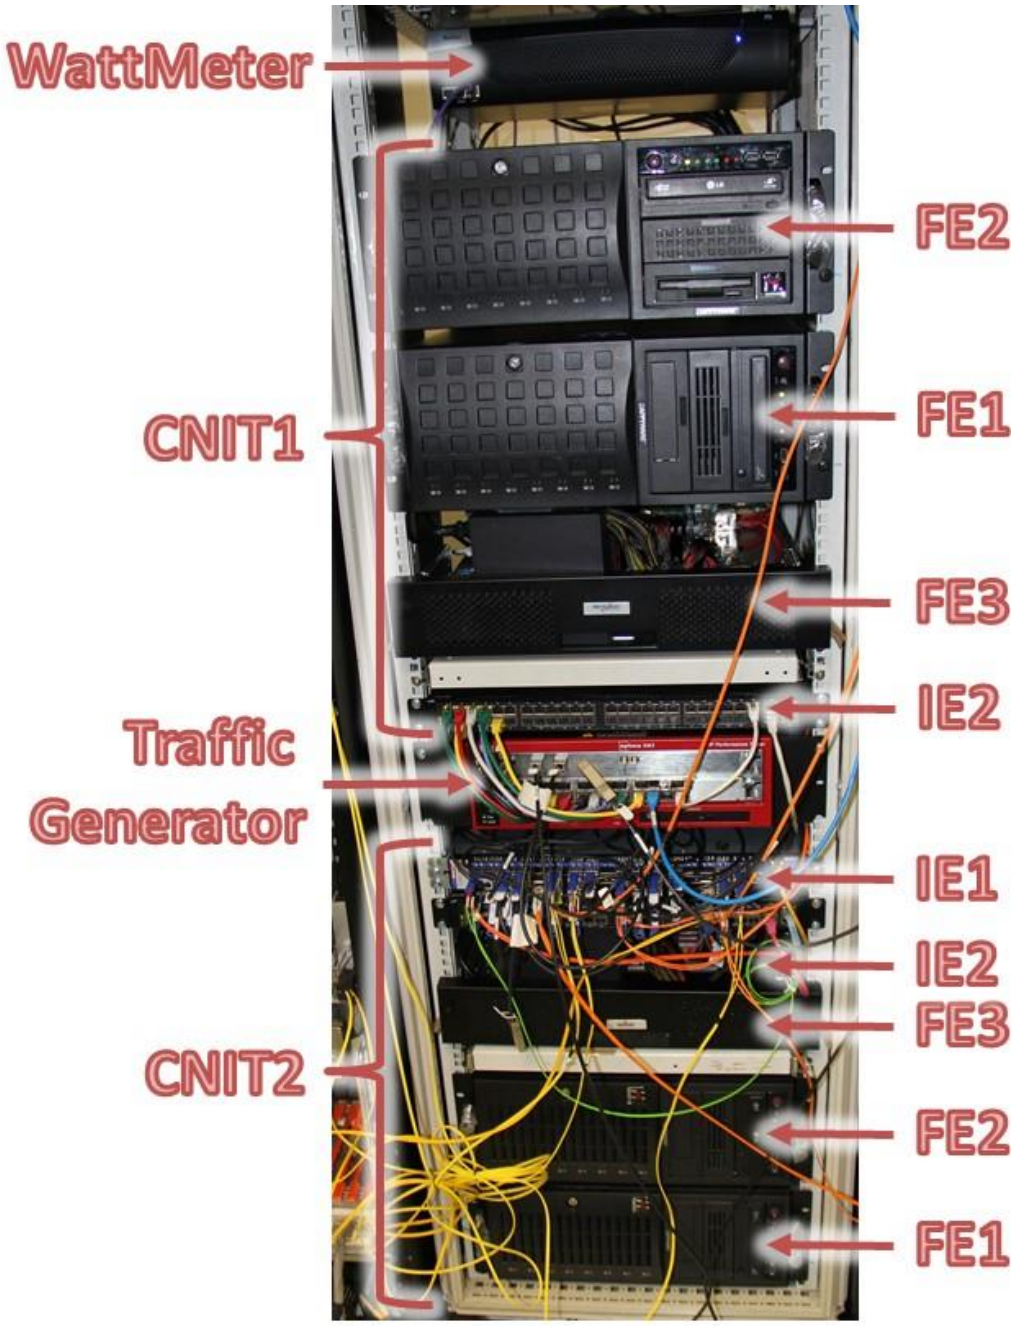

Fig. 6. The rack hosting the core network cloud and the Ixia traffic generator. The building blocks of the CNIT1 and CNIT2 instances of the DROP router are outlined.

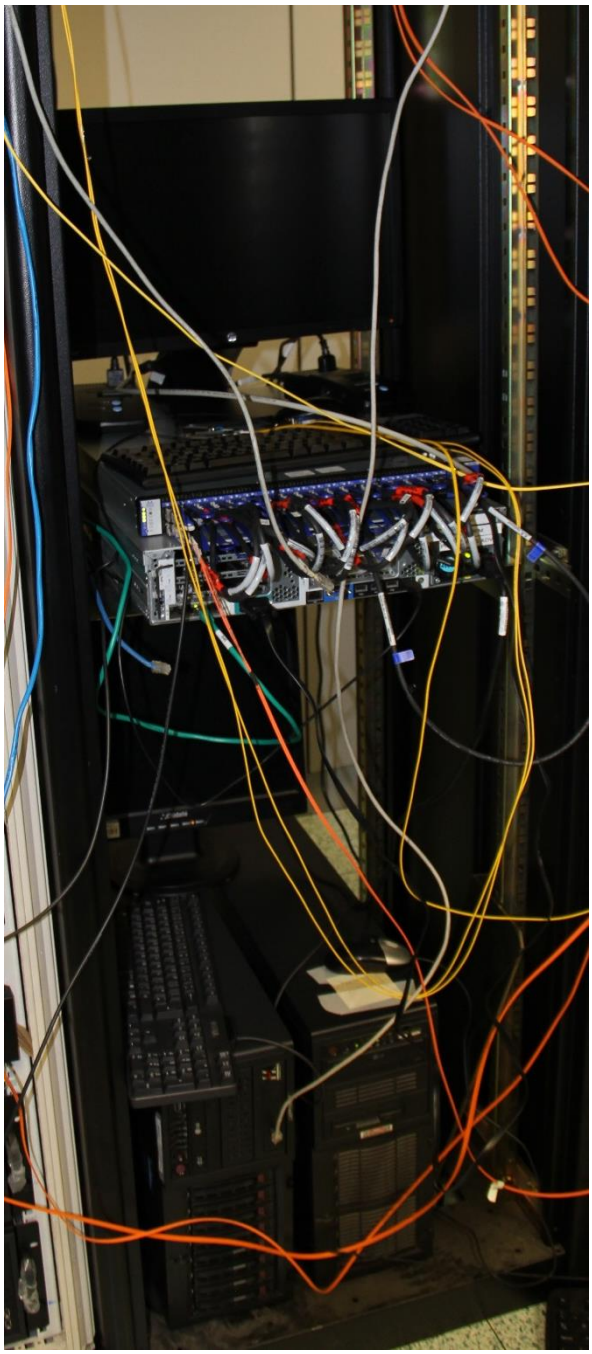

Fig. 7. The rack hosting the servers running the emulation networks (bottom part), and the datacentre (upper part).

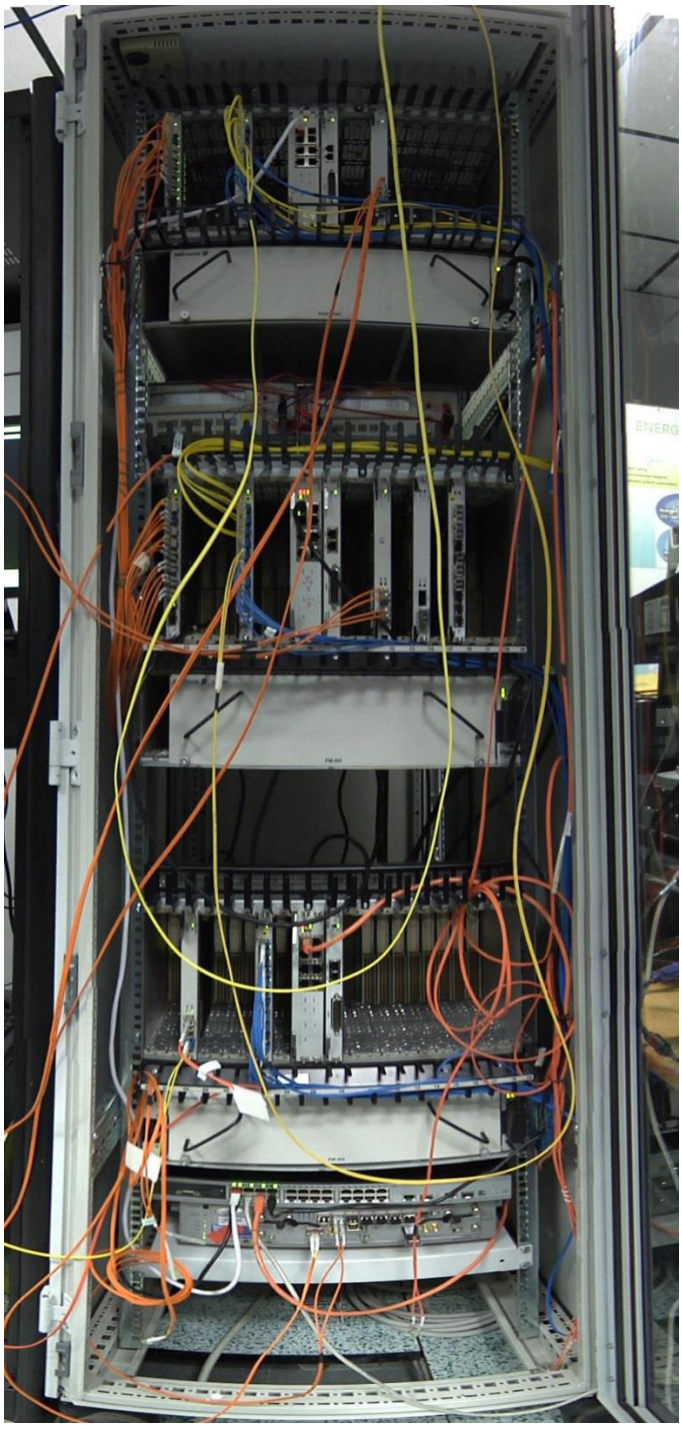

Fig. 8. The rack hosting the metro cloud.

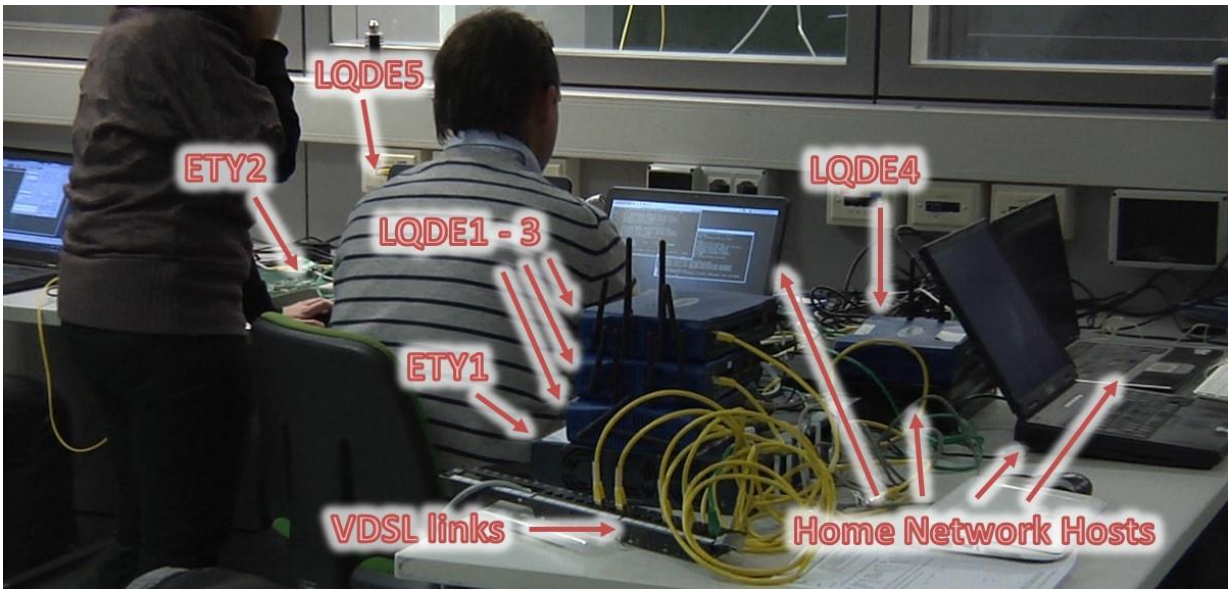

Fig. 9. The physical setup of the home and access clouds.

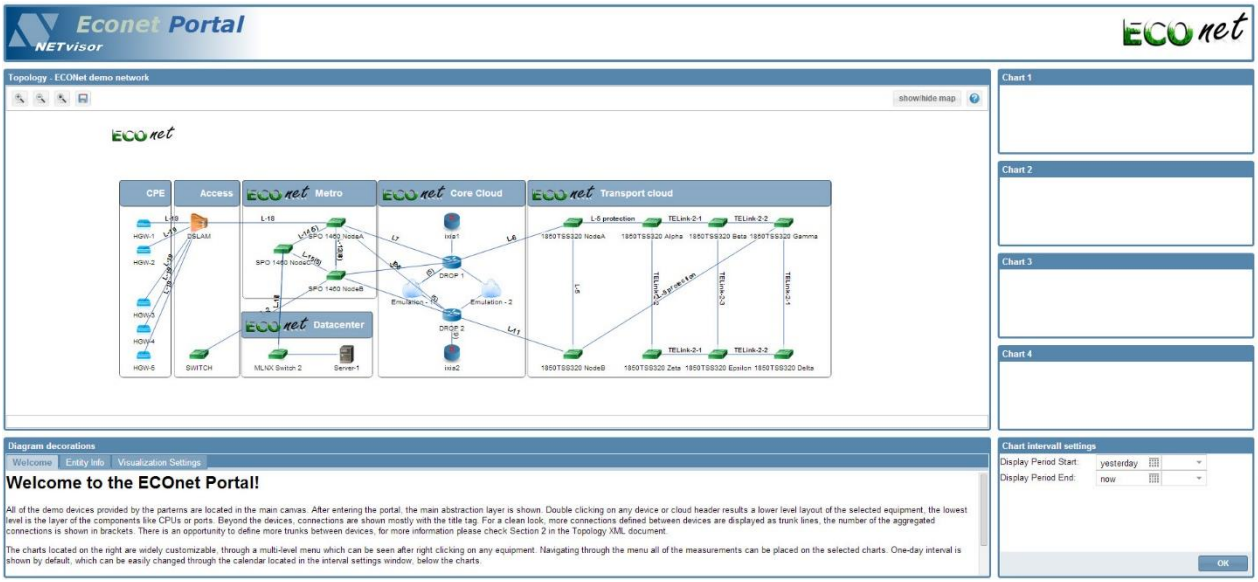

Fig. 10. Screenshot of the web GUI of the central energy-aware OAM system used in the demonstration.

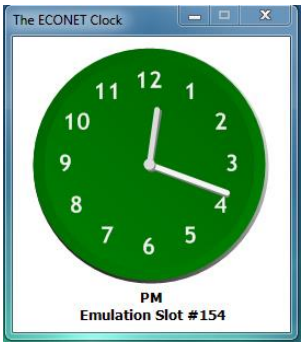

Fig. 11. Screenshot of the ECONET clock application.

The demonstrator was used in two main typologies of tests, which can be summarized as follows:

- Tests with traffic loads following daily patterns, which have been oriented to measure the impact of ECONET technologies at data- and control-planes on the energy consumption of the demonstration network;
- Tests to validate the performance and the functional behaviour of the Network Connectivity Proxy.

The two typologies of tests have been necessary, because of the different nature of the Network Connectivity Proxy with respect to all the other ECONET solutions evaluated in the demonstrator. On one side, the other ECONET solutions aim at extending network protocols and technologies for reducing the energy consumption of the network itself, by means of dynamic modulation of the trade-off between performance and power requirement. On the other side, the Network Connectivity Proxy directly targets the reduction of energy use in "networked" (or network-attached) hosts at the customer's home, which are generally considered "outside" the network.

Therefore, the Network Connectivity Proxy function and related validation/evaluation tests are fully complementary to the other parts integrated into the demonstrator. Having these considerations in mind, and for easing the description of the achieved results, the ECONET consortium decided to organize the experimental tests in the two typologies introduced above.

The following subsections are organized as follows. Subsections 3.1 and 3.2 report a summary of the experimental results obtained in the first typology of tests. In detail, Subsection 3.1 gives a global overview of the results obtained with the demonstrator as a whole, while Subsection 3.2 enters into the details of each single cloud, relating the global results of the previous section with the local behaviour of the prototypes in the cloud. Subsection 3.3 collects all the experimental results obtained for validating and evaluating the Network Connectivity Proxy.

3.1 Demonstrator network-wide results

The following demonstration network-wide results are reported and discussed in this section:

- The energy consumption values collected by the demonstrator prototypes, broken down per network cloud, in Fig. 12;
- The energy saving for each cloud, calculated against the Business-as-Usual (BAU) energy consumption, in Fig. 13;
- The energy efficiency expressed as "Gbps per W" per each network cloud according to the time slot in Fig. 14, and the increase of energy efficiency with respect to the BAU scenario in Fig. 15;
- The percentage of energy saving provided by the prototype with respect to their business-as-usual power requirements in Fig. 16;
- Throughput values of the traffic flows F-1,…, F-7 in Fig. 17;
- The transmission and reception loads of the demonstration network links in Figs. 18-41.

All these results are expressed according to the time slot number. It is worth recalling that the traffic flows are generated with a "night & day" profile divided into 300 time slots, each one representing approximately 5 minutes of a day. In each time slot, the offered load is maintained constant.

The results reported in this section and in the next one have been obtained by averaging the results of 10 test sessions. It is important to mention that, given the deterministic nature of traffic loads, the performed test sessions provided very stable measurements, and the difference between the results obtained in different test session was negligible (below 2% of any monitored performance indexes).

Preliminary tests have been performed in order to measure the business-as-usual energy requirement of the demonstrator. For this purpose, all the ECONET optimizations at data- and control-planes have been disabled on the prototypes, and the energy consumption has been measured with the same instrumentation used in the following test sessions. This energy consumption turned out to be approximately 2.6 kW, and it is almost constant according to the offered traffic load (a maximum variation of 10 W was observed).

Fig. 12 reports the raw energy consumption of the demonstrator according to the time slot, or, more precisely, the energy consumption measured on the physical network device prototypes in each cloud. This energy consumption has been reported in the form of stacked areas, and each area collects the overall energy absorption of the physical network prototypes belonging to it. The energy consumption estimated for the virtual nodes of the emulation networks in the core and the transport cloud is not considered in Fig. 12, but it will be later discussed in subsections 3.2.1 and 3.2.2, and their impact on the overall network energy requirement will be analysed in section 4.

In addition, the energy consumption coming from "networked" devices (i.e., the MLX server or the hosts placed in the home networks) is not part of the figure reported in Fig. 12. Excluding both emulated nodes and networked devices from Fig. 12 allows better highlighting the behaviour of physical prototypes, which integrate both data- and control-plane solutions as well as GAL implementations proposed in the previous ECONET WPs. Moreover, while the performance of virtual nodes can be only estimated (and it has been done with high-accuracy models of the network devices), the behaviour of physical prototypes can be directly measured by third-party professional measurement instrumentation (e.g., traffic generators, watt-metres, etc.).

Therefore, the results in Fig. 12 and Fig. 13, which outline the savings that are achieved with respect to the BAU case, represent a solid proof of how the overall reference architecture (i.e., green data-plane + GAL + green control-plane) and the specific ECONET solutions can jointly work in a complex testing scenario (with many cross-technology interoperability points), by providing significant reduction of network power requirements.

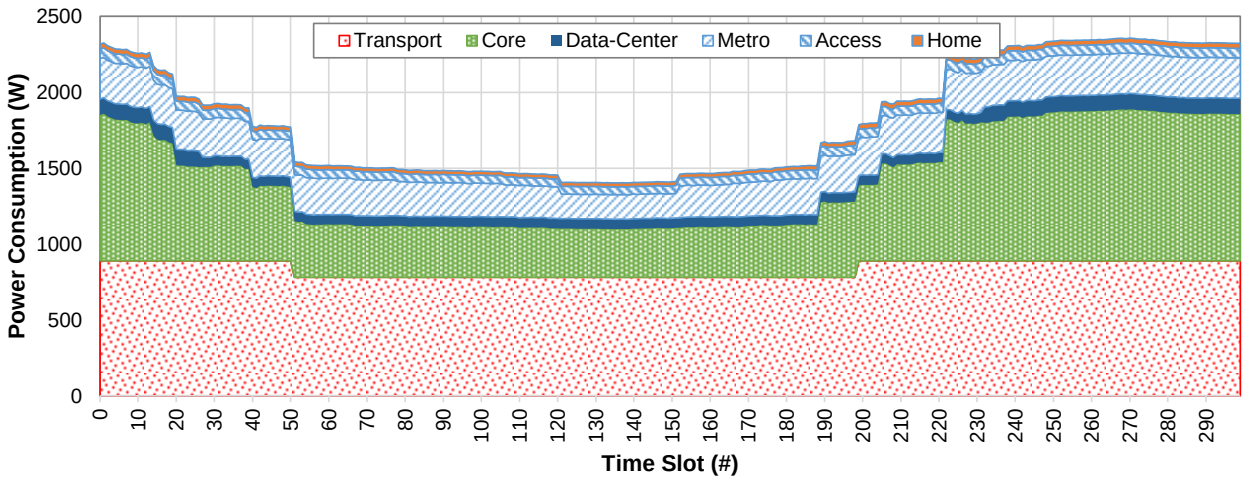

Fig. 12. Energy consumption of the demonstrator clouds (excluded the two emulation networks in the core cloud) according to time slots. The results are represented as stacked areas in order to make evident the overall demonstrator energy consumption.

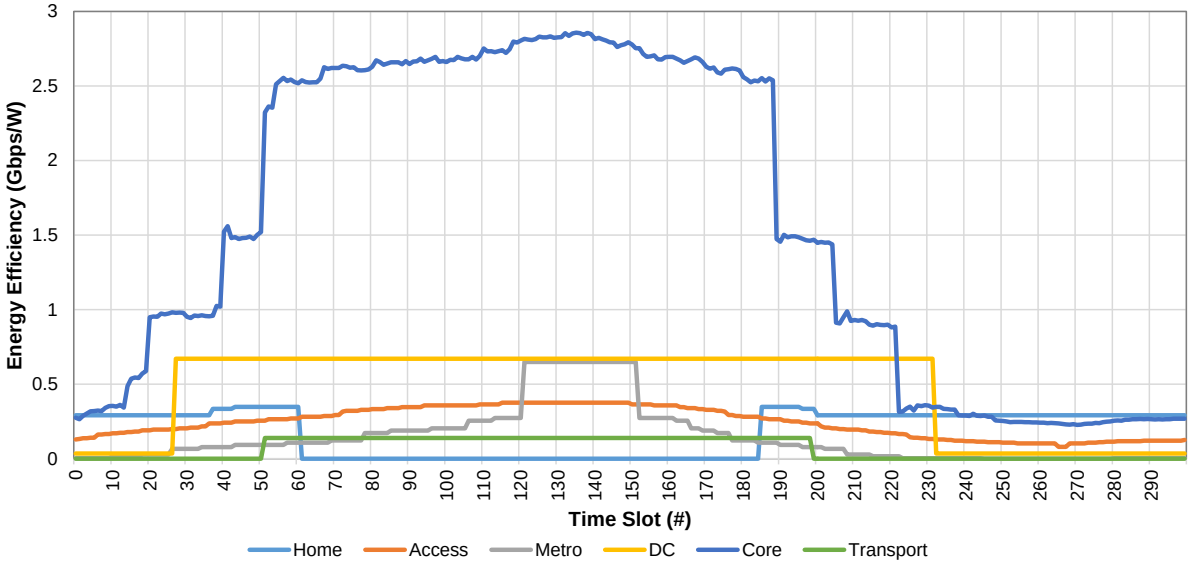

Fig. 13. Energy saving of the demonstrator clouds (excluded the two emulation networks in the core cloud) according to time slots and with respect to the Business-as-Usual energy consumption.

The overall energy consumption of the demonstrator passes from a maximum power consumption of approximately 2.35 kW during the "day" slots to a minimum absorption of about 1.3kW during the night slots. The consumption profile appears to closely follow the shape of "night & day" traffic fluctuations, demonstrating the effectiveness of the green mechanisms included in the prototypes.

Observing Fig. 12 it is also evident how different design and integration choices have been undertaken in the different clouds. For instance, the energy savings provided by the transport cloud result sensibly lower than the ones of the other clouds, since, in this specific case, the energy modulation is not done on the actual traffic load, but based on the maximum bandwidth reserved to each flow during the night and day hours. This design choice is fully in line with bandwidth reservation approaches usually adopted in transport networks. The clear effect is the impossibility for this cloud to closely follow the dynamics of traffic load, and a consequent over-dimensioning of the resources left active in each time slot. Additional details in this respect can be found in subsection 3.2.2.

The core cloud is the second absolute contributor (after the transport cloud) to the overall energy consumption of the demonstrator, but it has been dimensioned to carry volumes of traffic much larger than the other demonstration clouds. The two instances of the DROP router need to achieve an aggregated throughput of more than 50 Gbps. So, despite the facts that:

(i) the DROP instances are prototypes closer to the research world, rather than to a commercial product,

(ii) the DROP internal architecture is based on the Network Functions Virtualization (NFV) paradigm, which surely extends network programmability at the cost of using more energy-hungry hardware platforms,

the energy per bit efficiency provided by this prototype (see Figs. 14 and 15) is (surprisingly) comparable with the one of commercial prototypes. Additional details in this respect can be found in subsection 3.2.2.

The energy-efficiency levels of the clouds, in terms of carrying Gbps / W shown in Fig. 14, provide additional interesting information on the high-level behaviour of the demonstrator. The energy efficiency level for each cloud has been calculated by summing all the flows entering a cloud for the number of times that they enter, and dividing the result by the overall energy consumption of the same cloud. For instance, considering the datacentre cloud, the traffic volume considered was twice the throughput of the F-3, F-5, and F-7 flows, since these flows enter from the L-16, L-17 and the injection point 4, and are subsequently re-injected back from the MLX server. The energy consumption considered in this calculation is only the one of the MLX switch. Similar calculations have been performed for the other clouds.

It must be noted that these results are strictly related to the cloud architecture and traffic flow definitions, and, consequently, they cannot be interpreted as a characterization of the prototypes (which rather has to be performed by isolating single clouds/devices and filling them with suitable levels of traffic).

Notwithstanding these considerations, we can observe how the energy modulation mechanisms integrated in the prototypes allows maintaining a high level of efficiency also during night hours, when traffic volumes are low. In this respect, Fig. 15 reports the difference between the values of Fig. 14 and the ones obtained in the BAU scenario, and highlights that ECONET technologies allow energy efficiency increases by factors of 50-70% during night hours with respect to the BAU scenarios.

From a general perspective, Figs. 14 and 15 confirm that devices and network segments working near the core can provide better energy efficiency levels, since each device works with large volumes of aggregated traffic. It can be noted that these levels of efficiency depend on the internal topologies of the various clouds: where clouds are composed of multiple devices, the efficiency degree tends to be lower than in the case with single devices. This effect is clearly caused by the fact that internal interconnections are not considered in the energy-efficiency level calculation.

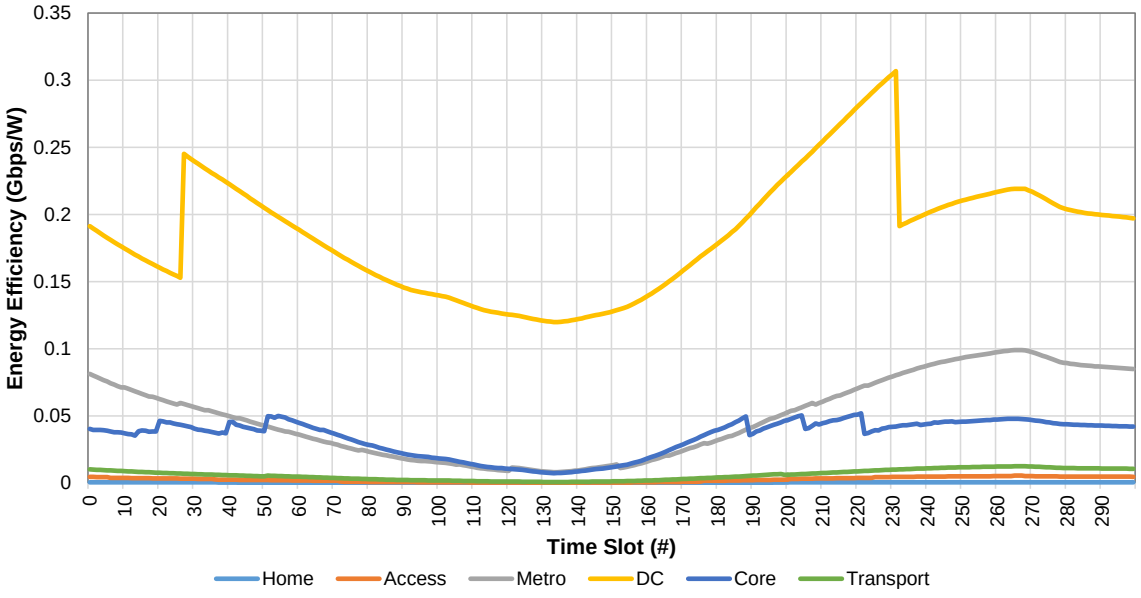

Fig. 14. Energy efficiency in terms of Gbps per W according to the time slot during the experimentation of the ECONET technologies. The traffic volumes considered are the F-1,…, F-7 weighted for the number of times that they enter a cloud.

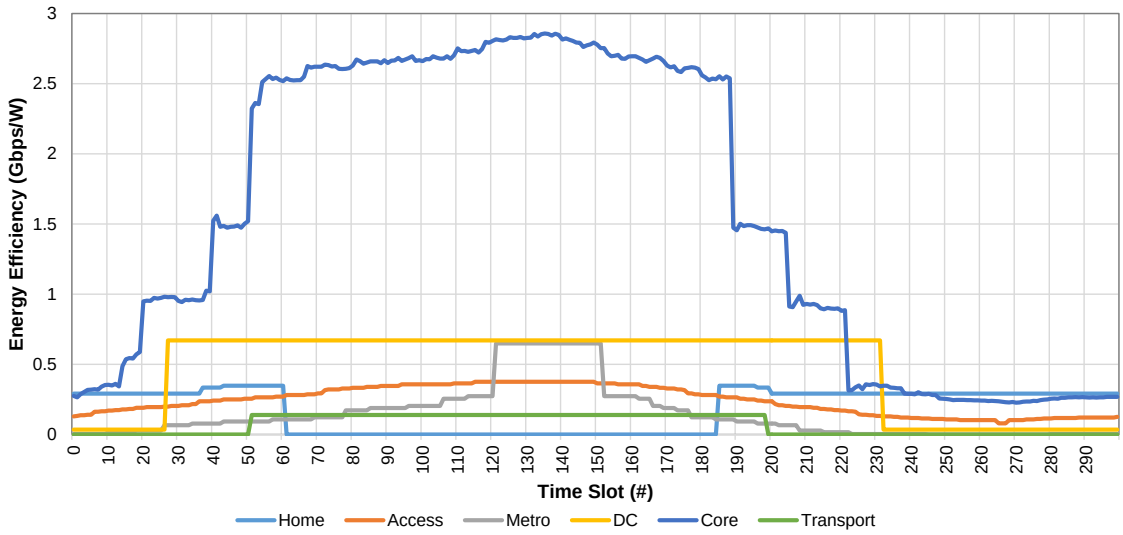

Fig. 15. Increase of energy efficiency in terms of Gbps per W according to the time slot provided by the ECONET technologies with respect to the BAU scenario.

Fig. 16, which reports the energy saving with respect to the business-as-usual case according to the time slot, makes evident that the green mechanisms allow decreasing power requirements by a factor of 50% during night hours, while in rush hours the savings are approximately 10% of the business-as-usual energy absorption.

However, it is worth to underline how the pictures represented in Figs. 12 and 16 are not representative of the impact of the ECONET technologies in real network infrastructures, given a number of aspects that have to be carefully taken into account. Among the most relevant aspects, we can cite:

- For obvious reasons, the demonstrator cannot provide *typical densities of devices* usually deployed in network segments of Telecom Operator infrastructures. For instance, the demonstrator includes 2 core routers for only 2 access devices and 5 home gateways. Therefore, the results coming from the demonstrator clouds need to be suitably weighted in order to represent typical deployments of the different network segments.
- As mentioned in the D6.4 report (Sect. 2) [1], the ECONET prototypes have been completed with *"legacy" components* (without green mechanisms) in order to provide the interconnection and the interoperability among clouds and to realize topologies suitable to test the control-plane algorithms and protocols proposed in the project.
- *The dimensioning of the links and of the devices with respect to the traffic load* is not in line with the one used in real network deployment. In many clouds, the traffic loads are sensibly larger than in real networks, since they exceed half of the deployed (link/device) capacities. This over-dimensioning of traffic loads, which obviously causes lower energy saving, has been decided for correctly triggering the energy saving mechanisms in all the clouds (and, then, also

in the non-overloaded ones – e.g., the datacenter cloud), and for providing a worst-case bound of the potential impact of the project.

- Some of the *prototypes do not include all the possible energy-saving solutions* studied in the project, since, as foreseen in the project scheduling and in the proposal itself, hardware/software integration and development activities to include all these mechanisms go beyond the dimension and the scope of the project. It is worth recalling that the demonstrator aims at giving a proof-of-concept indication of the effectiveness of the ECONET technologies, and that it is not, and it must not be, a pre-commercial release of products.

Beyond these aspects, the ECONET Consortium believes that the raw energy saving obtained constitutes not only a precious feedback on the effectiveness of the single solutions, but rather demonstrates how these technologies can have a disruptive impact on the network energy consumption, notwithstanding their early prototyping status and the presence of a significant share of legacy components.

Having all these points in mind, the remainder of this section will focus on the demonstration results in order to highlight the behaviour and the effect of the various solutions integrated into the demonstrator. Later in section 4, the experimental results will be re-analysed and re-elaborated in order to give an accurate estimate of the potential impact of the ECONET technologies when massively applied to a real infrastructure of network operators.

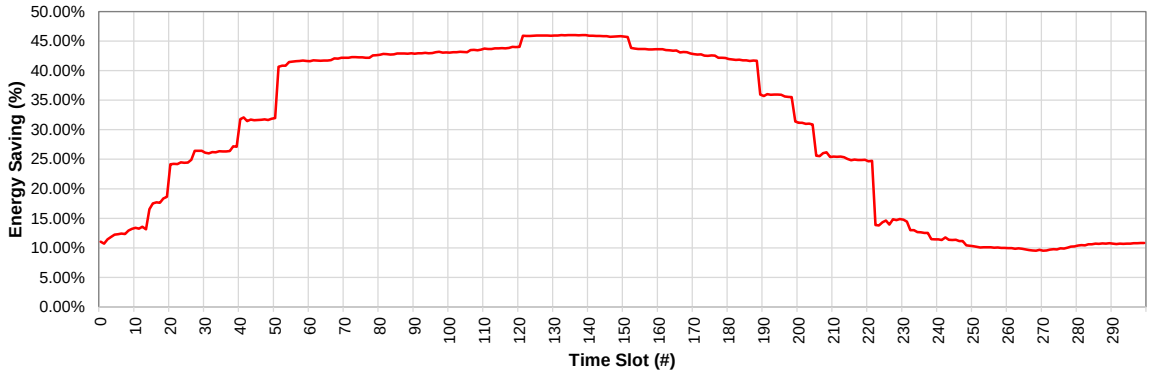

Fig. 16. Energy savings obtained with the demonstrator (excluded the two emulation networks in the core cloud) with respect to the BAU energy consumptions. The values are expressed according to time slots.

Coming back to the analysis of demonstration results, it can be noted from Fig. 17 how the demonstration network exhibited throughput values for the F-1, …, F-7 in line with their offered load. In fact, testing instrumentation did not report any packet losses during experimentation activities. The overall network throughput passes from a maximum value of 30 Gbps (near slot #270) to a minimum value of approximately 6 Gbps (near slot #135).

The bandwidth occupancy of almost all the demonstrator links is reported in Figs. 18-41. As expected, the bandwidth occupancy values of links follow very close the fluctuation of the F-1,…, F-7 flows, since they come from the superposition of the throughputs of those flows that are crossing them (see Figs. 2 and 3 for a high-level overview of the flow paths in the demonstration network).

Despite the general behaviour of link occupancies, it is interesting to note the effects of network control policies. In fact, as evident from the steps in the occupancy of a number of links (e.g., L-4, L-6, L-10, L-11, etc.), optimization algorithms acting in the various clouds decide, when the traffic load

is sufficiently low, to re-route the traffic for shutting down link interfaces and/or internal components of network nodes.

In more detail, the following re-routing events occur:

- At slot #51, F-4 is re-routed from L-4 and L-10 to L-6 and L-11. The re-routing is decided by the CNIT2 router, and allows the transport network to pass from the day configuration to the night one (see Figs. 20, 22, 26 and 27);
- At slot #200, F-4 is re-routed back to L-4 and L-10 (this time from L-6 and L-11). The re-routing is decided by the CNIT2 router, and allows the transport network to pass from the night configuration to the day one (see Figs. 20, 22, 26 and 27);
- At slot #78, F-6 is re-routed from L-12 to L-13 by the green PCE acting in the metro cloud (see Figs. 28 and 29), and L-12 is put into standby mode.
- At slot #104, F-6 is re-routed from L-13 to the path L-14 + L-15 by the green PCE acting in the metro cloud (see Figs. 29, 30 and 31).
- At slot #112, F-4 is re-routed from L-13 to the path L-14 + L-15 by the green PCE acting in the metro cloud (see Figs. 29, 30 and 31), and L-13 is put into standby mode.
- At slot #165, L-13 is waking up, and F-4 is re-routed back to it from the path L-14 + L-15 (see Figs. 29, 30 and 31).
- At slot #170, also F-6 is re-routed back L-13 from the path L-14 + L-15 (see Figs. 29, 30 and 31).
- At slot #184, L-12 is waking up, and F-6 is re-routed back to it from L-13 (see Figs. 28 and 29).

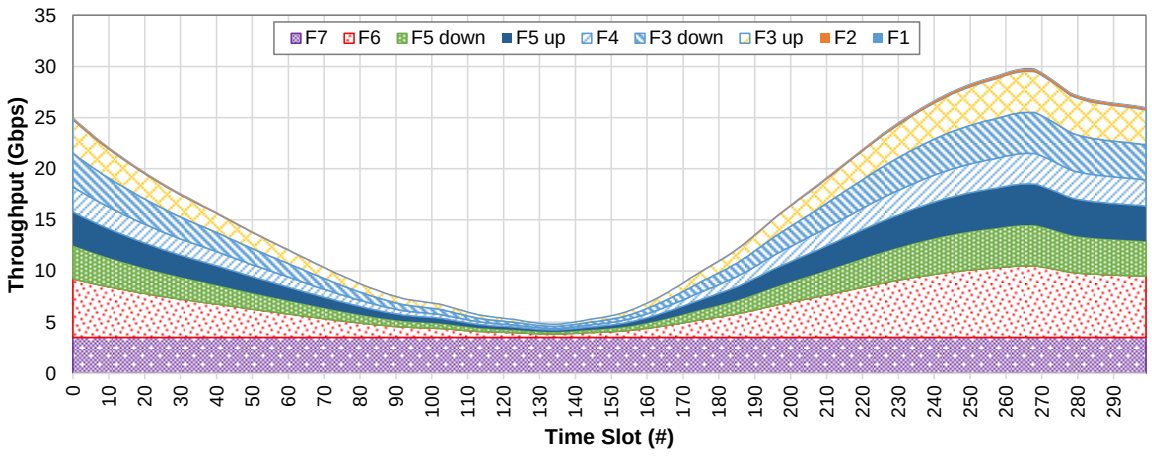

Fig. 17. Throughput of the traffic flows according to time slots. The results are represented as stacked areas in order to highlight the overall network throughput.

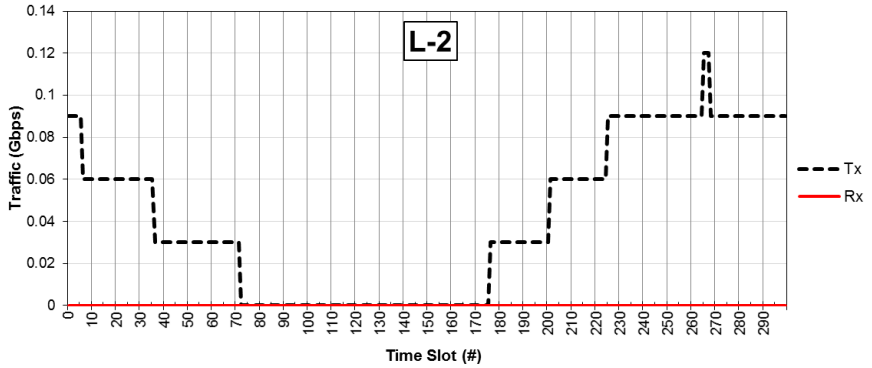

Fig. 18. Transmission and reception bandwidth utilizations according to time slots on the L-2 link.

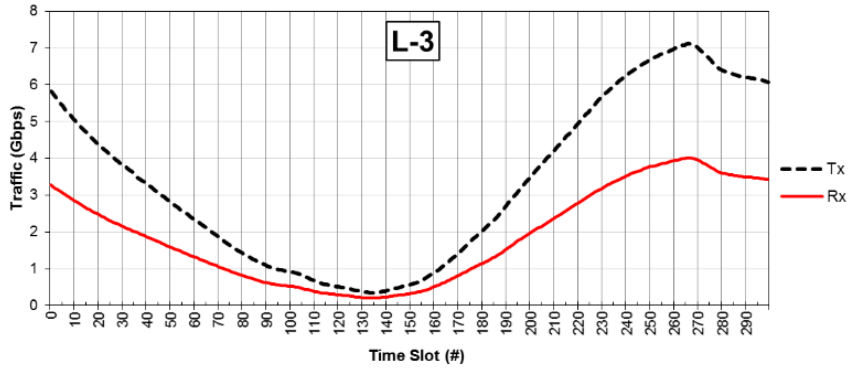

Fig. 19. Transmission and reception bandwidth utilizations according to time slots on the L-3 link.

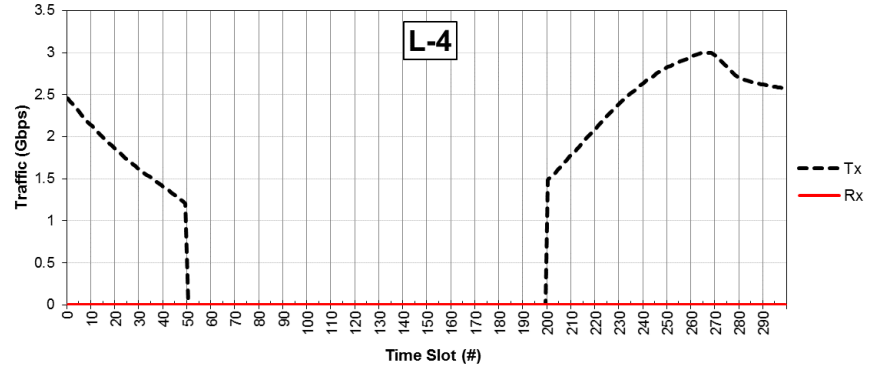

Fig. 20. Transmission and reception bandwidth utilizations according to time slots on the L-4 link.

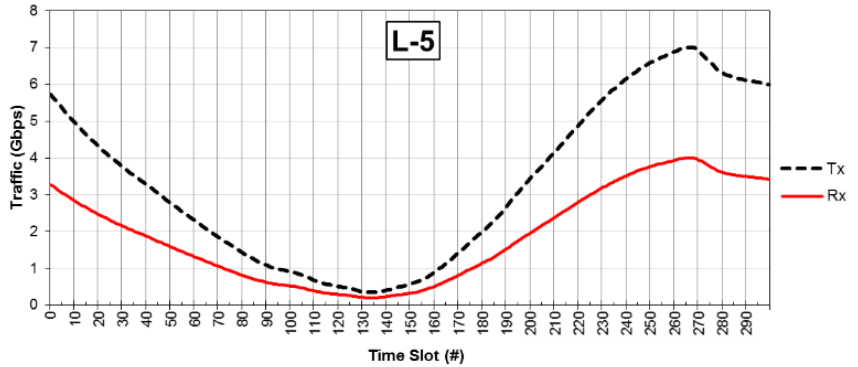

Fig. 21. Transmission and reception bandwidth utilizations according to time slots on the L-5 link.

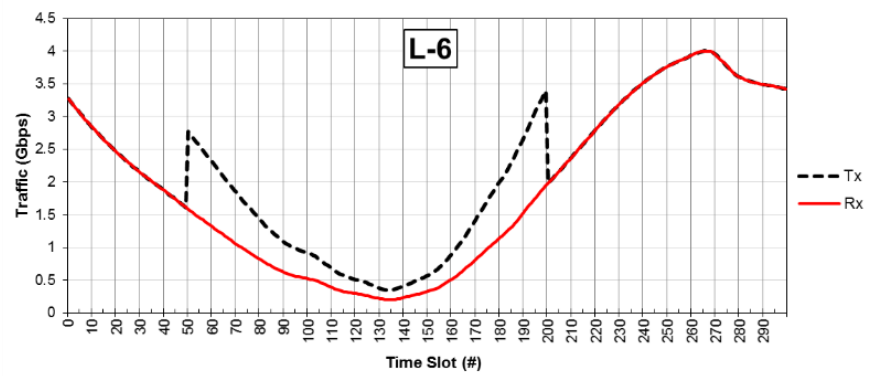

Fig. 22. Transmission and reception bandwidth utilizations according to time slots on the L-6 link.

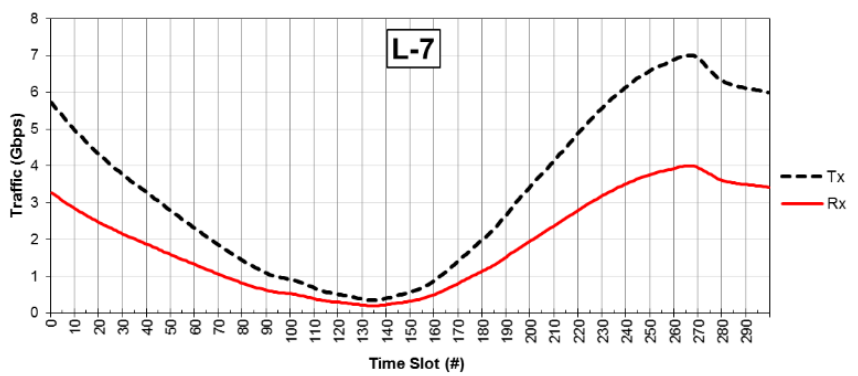

Fig. 23. Transmission and reception bandwidth utilizations according to time slots on the L-7 link.

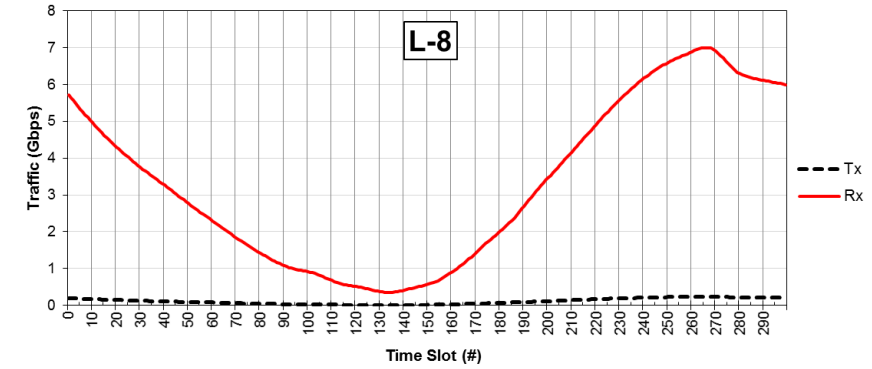

Fig. 24. Transmission and reception bandwidth utilizations according to time slots on the L-8 link.

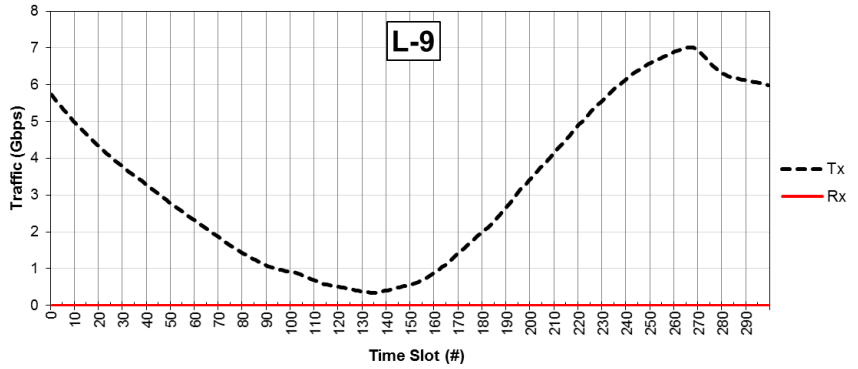

Fig. 25. Transmission and reception bandwidth utilizations according to time slots on the L-9 link.

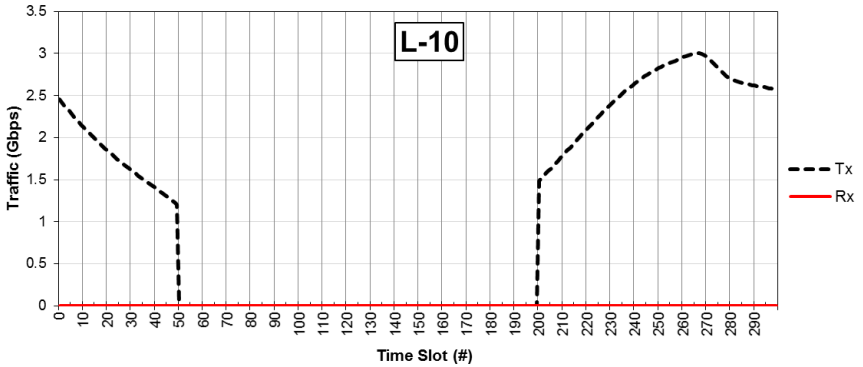

Fig. 26. Transmission and reception bandwidth utilizations according to time slots on the L-10 link.

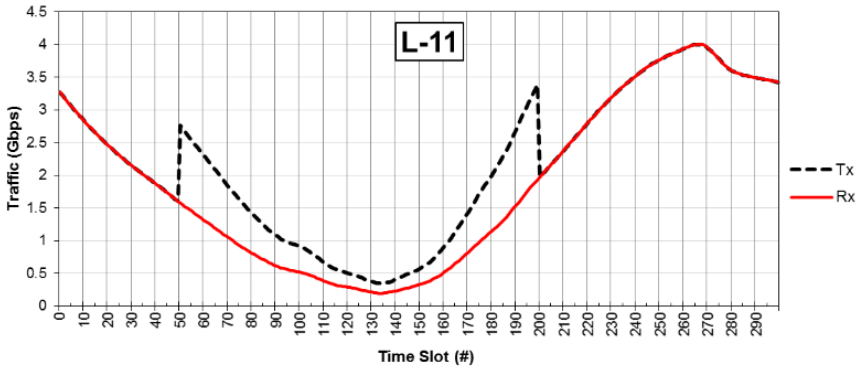

Fig. 27. Transmission and reception bandwidth utilizations according to time slots on the L-11 link.

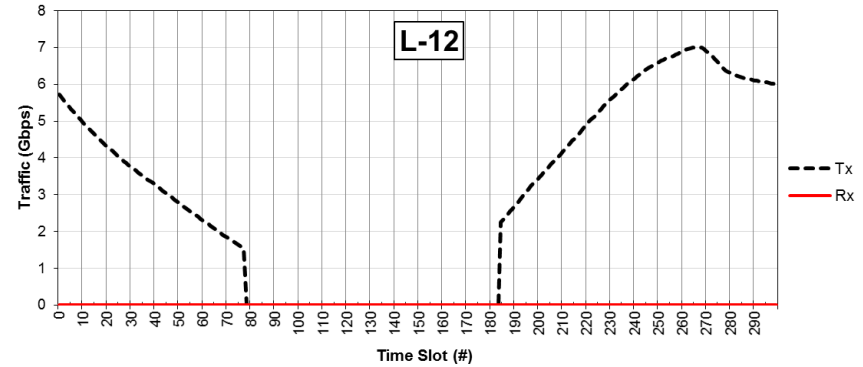

Fig. 28. Transmission and reception bandwidth utilizations according to time slots on the L-12 link.

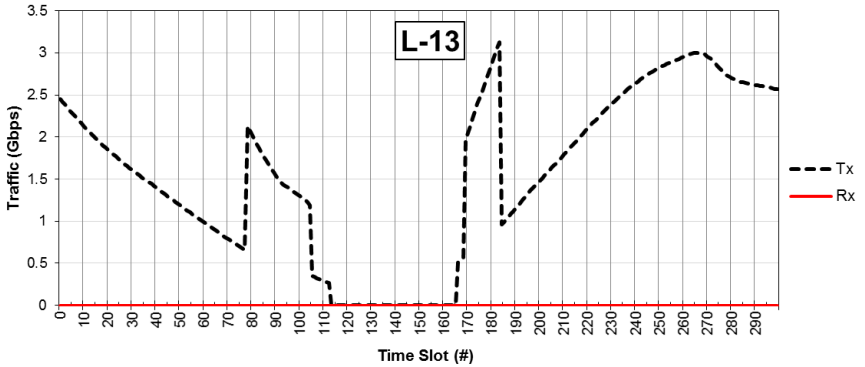

Fig. 29. Transmission and reception bandwidth utilizations according to time slots on the L-13 link.

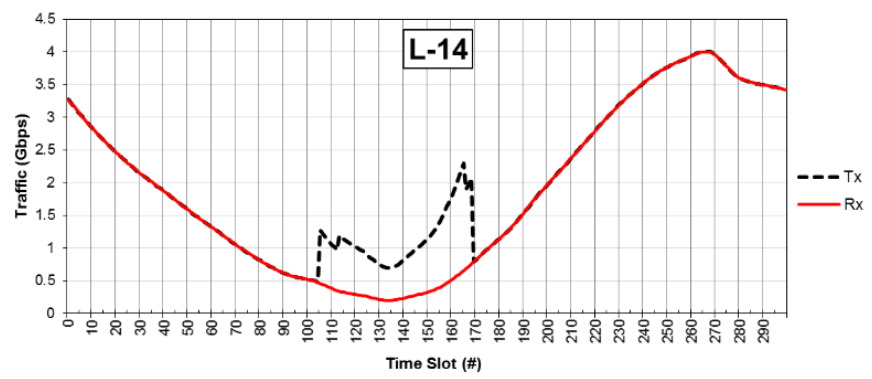

Fig. 30. Transmission and reception bandwidth utilizations according to time slots on the L-14 link.

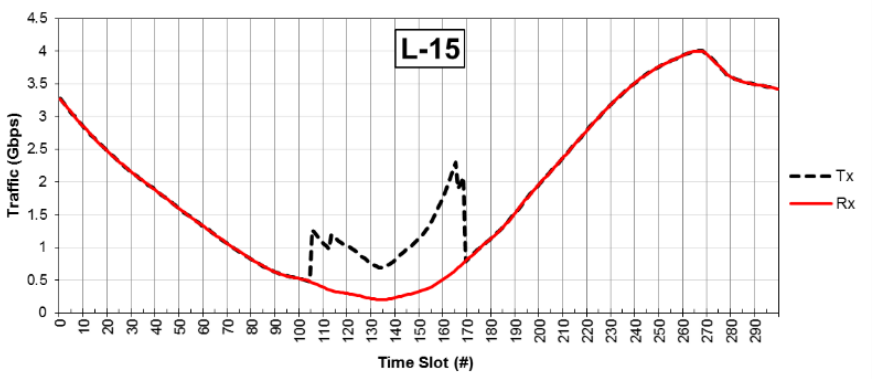

Fig. 31. Transmission and reception bandwidth utilizations according to time slots on the L-15 link.

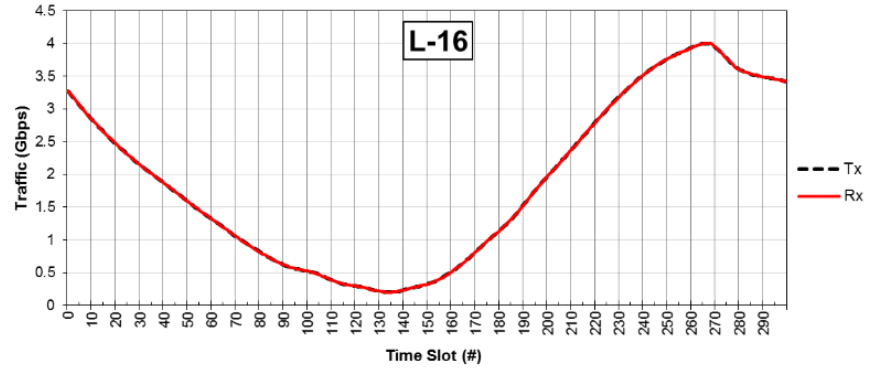

Fig. 32. Transmission and reception bandwidth utilizations according to time slots on the L-16 link.

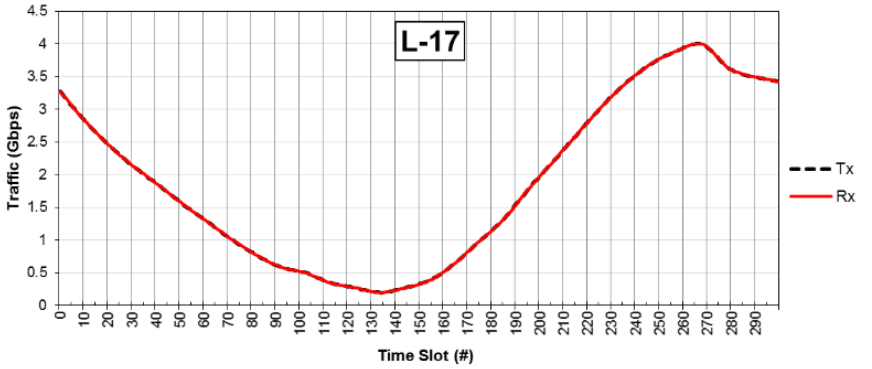

Fig. 33. Transmission and reception bandwidth utilizations according to time slots on the L-17 link.

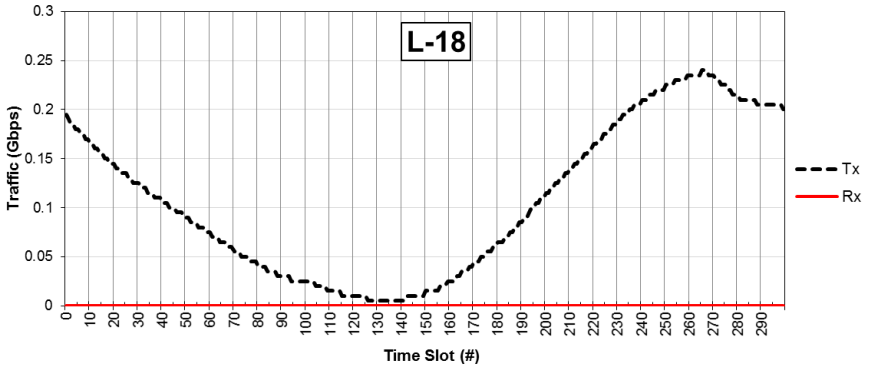

Fig. 34. Transmission and reception bandwidth utilizations according to time slots on the L-18 link.

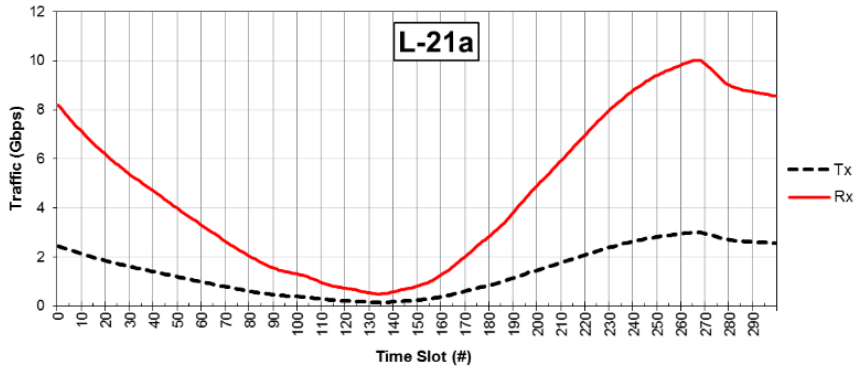

Fig. 35. Transmission and reception bandwidth utilizations according to time slots on the L-21a link.

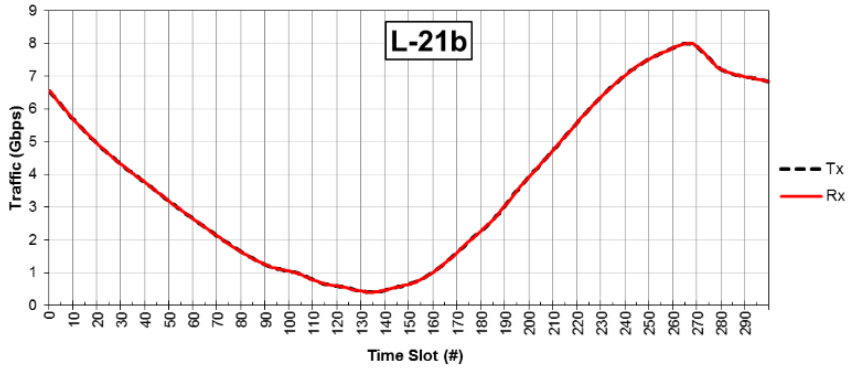

Fig. 36. Transmission and reception bandwidth utilizations according to time slots on the L-21b link.

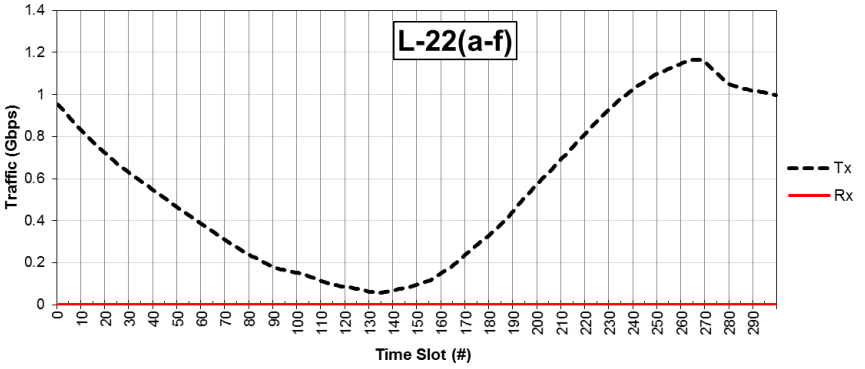

Fig. 37. Cumulative transmission and reception bandwidth utilizations according to time slots on the L-22(a-f) links.

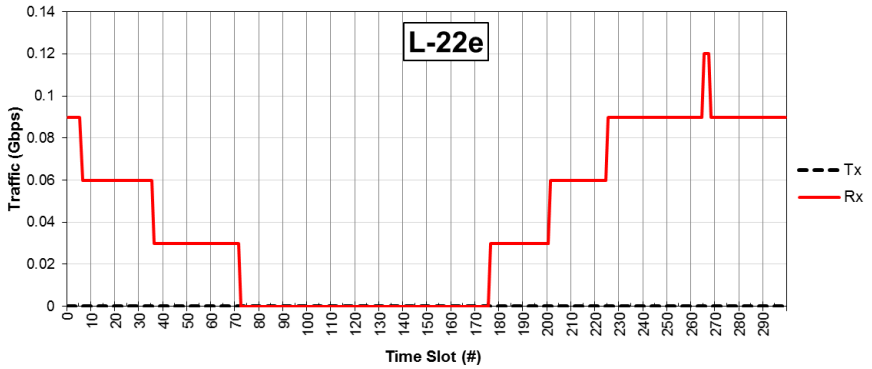

Fig. 38. Transmission and reception bandwidth utilizations according to time slots on the L-22e link.

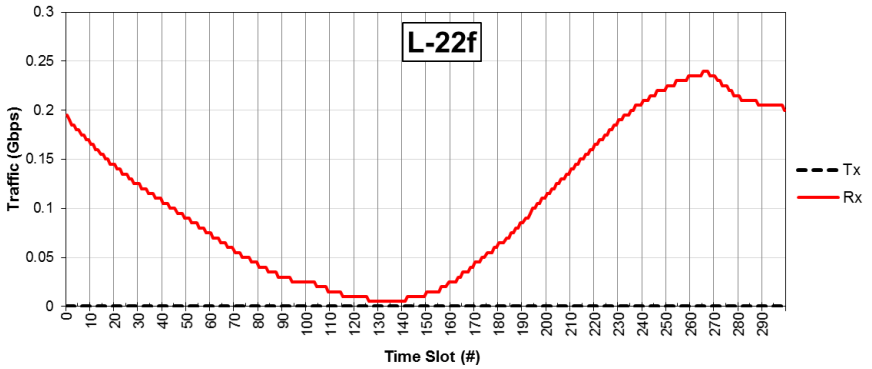

Fig. 39. Transmission and reception bandwidth utilizations according to time slots on the L-22f link.

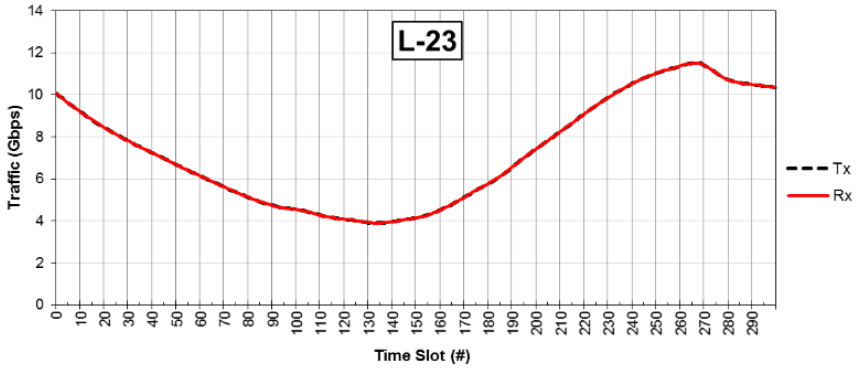

Fig. 40. Transmission and reception bandwidth utilizations according to time slots on the L-23 link.

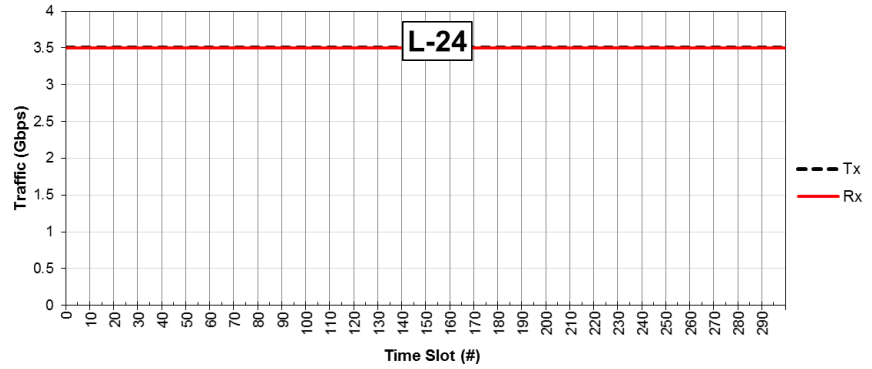

Fig. 41. Transmission and reception bandwidth utilizations according to time slots on the L-24 link.

In order to have a complete picture of the demonstrator performance and of the impact of ECONET technologies on the Quality of Experience (QoE) perceived by end-users, additional end-to-end measurements have been performed. In detail, the TCP test application [2] was used to download a file of 14 MB from the MLX server in the datacenter cloud to a PC connected to the LQDE-4 HG in the home cloud. The tests have been performed only every 10 slots, and in the presence of the traffic and of the power configurations described above. Each download test in a specific slot has been repeated 10 times.

Fig. 44 reports the obtained results in terms of the average download time of the file, and of maximum and minimum deviation from the average. As can be noted, despite the fluctuation of traffic and the large number of power management mechanisms acting at data- and control-planes, no performance variations are visible across all the examined time slots. Therefore, Fig. 44 clearly indicates that a reasonable adoption and activation of ECONET technologies allows saving a significant amount of energy without affecting the QoE level perceived by end-users.

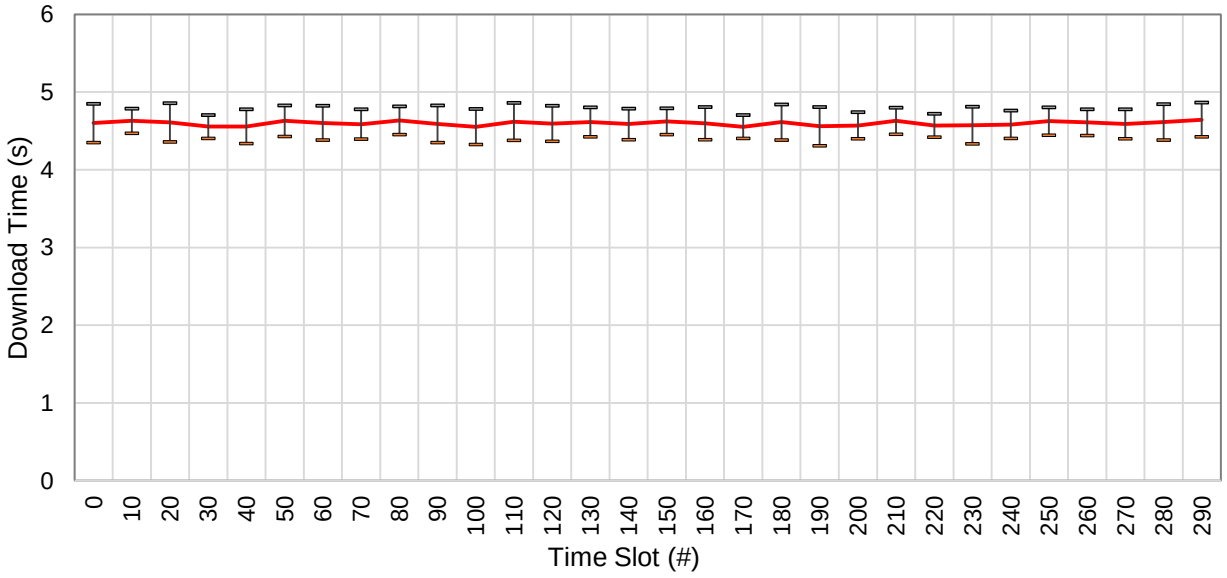

Fig. 42. Download time according to the time slot of the TCP connection used to evaluate the user's Quality of Experience. The chart reports the average value and the upper and lower bounds of the obtained measurements.

3.2 Focusing on single clouds

This section reports detailed results for each demonstration cloud obtained in the main experimentation, or collected in minor tests performed to better evaluate the behavior of specific technologies or solutions. The remainder of this section is organized as follows. Section 3.2.1 discusses the results obtained in the transport cloud. Section 3.2.2 focuses on the core cloud. The metro cloud is analyzed in section 3.2.3. The Datacenter cloud is in section 3.2.4. Finally, sections 3.2.5 and 3.2.6 report the results obtained in the access and home clouds, respectively.

3.2.1 Transport Network Cloud

As already reported (in the D6.4 report [1] – mainly in sections 2 and 6) the most relevant contribution of power saving in the Optical Transport Network is based on the ability to switch-off the unused resources. The transport network cloud used in the ECONET demonstrator shows this capability by using a setup based on two real nodes and six emulated nodes. Power consumption estimation is based on actual real-time measurements for the real nodes, while a model based on off-line measurements is used for the emulated network.

According to the general strategy of the ECONET demonstrator, this network segment follows a daily traffic pattern characterized by a peak load variation from day (20 Gbps) to night (10 Gbps).

This traffic variation is managed in a network where some resources are pre-allocated to background traffic to give the feeling of how the power management capability can actually run. It is clear that the power saving depends on the network topology, on the network size and on the capability to capture significant traffic load variations. This is to say that the measured saving cannot represent the general case, although it can give a feeling of the possible results.

Fig. 43 depicts the network topology and the traffic status designed for the ECONET demonstrator.

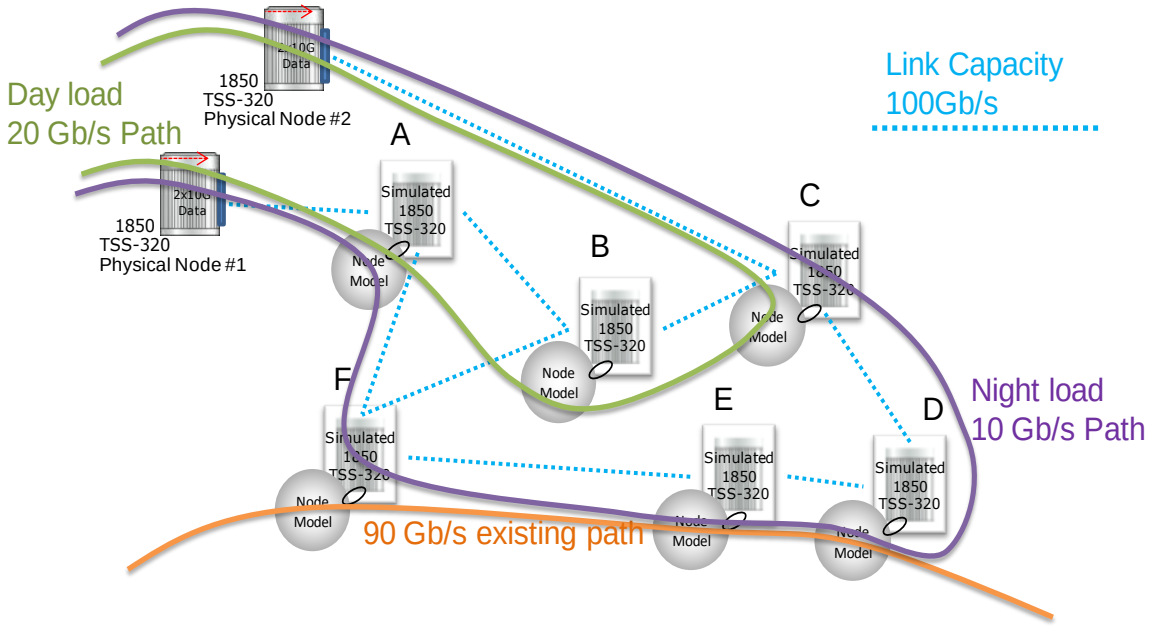

Fig. 43. Transport topology (physical and virtual nodes) and emulated flows. The green and the violet lines correspond to the paths of the F-4 and F-3 flows in the transport cloud during the day and night configurations, respectively.

The Night load can be forwarded along nodes A-F-E-D-C, being able to switch off node B, while the Day load exceeds the link capacity and has to be forwarded along the A-B-C path. During the demo the following results have been captured:

- without power management: Total power Day = 5262 W, Total power Night = 5242 W, Delta=0.4%
- with power management: Total power Day = 5262 W, Total power Night = 4693 W, Delta=11%

Using ECONET technology, implementing Day/Night path switching allows around 10.5% of saving (see Fig. 44).

Tables I-VIII provide the breakdown of the power consumption for each sub-component of the physical and virtual nodes.

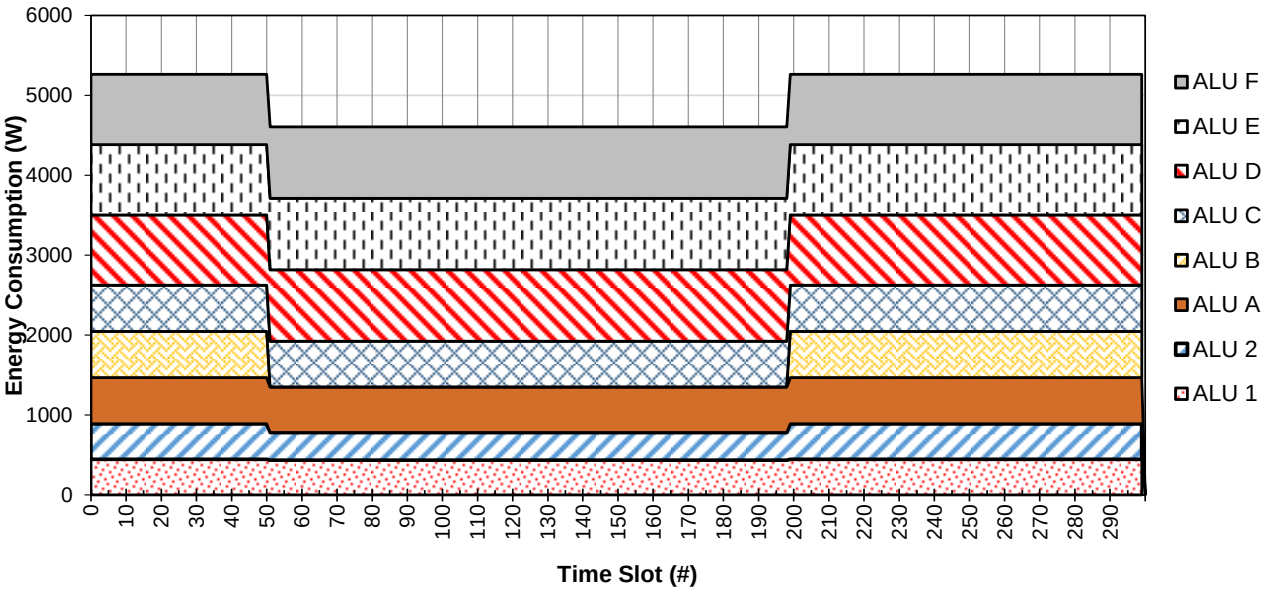

Fig. 44. Cumulative energy consumption of all the nodes in the transport cloud (both physical and emulated ones). The results are visualized in terms of stacked areas in order to express the total energy consumption of the transport cloud.

Table I. Breakdown of energy consumed by the physical node ALU1 in both the day and night configurations.

						DAY			Night			Saving
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #1		TSS320				1		**445.0**	1		**346.5**	**22.1%**
Common Parts							40	250.5		20	249.3	
Common Part #	1	FAN	320.0	30.0	0.0	1	40	30.0	1	20	30.0	
Common Part #	2	Supply	320.0	50.0	0.0	1	40	50.0	1	20	50.0	
Common Part #	3	EC	320.0	20.0	0.0	1	40	20.0	1	20	20.0	
Common Part #	4	SC	320.0	18.0	0.0	1	40	18.0	1	20	18.0	
Common Part #	5	Matrix 1	320.0	130.0	20.0	1	40	132.5	1	20	131.3	
Common Part #	6	Matrix 2	320.0	130.0	20.0	0	40	0.0	0	20	0.0	
Line Card #	1	2x10Gb/s	20.0	88.0	6.3	1	20	97.3	1	20	97.3	
Interface #	1	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	1	10	1.5	
Interface #	2	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	1	10	1.5	
Line Card #	2	2x10Gb/s	20.0	88.0	6.3	1	20	97.3	0	0	0.0	
Interface #	1	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	0	0	0.0	
Interface #	2	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	0	0	0.0	

Table II. Breakdown of energy consumed by the physical node ALU2 in both the day and night configurations.

						DAY			Night			Saving
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #1		TSS320				1		**445.0**	1		**346.5**	**22.1%**
Common Parts							40	250.5		20	249.3	
Common Part #	1	FAN	320.0	30.0	0.0	1	40	30.0	1	20	30.0	
Common Part #	2	Supply	320.0	50.0	0.0	1	40	50.0	1	20	50.0	
Common Part #	3	EC	320.0	20.0	0.0	1	40	20.0	1	20	20.0	
Common Part #	4	SC	320.0	18.0	0.0	1	40	18.0	1	20	18.0	
Common Part #	5	Matrix 1	320.0	130.0	20.0	1	40	132.5	1	20	131.3	
Common Part #	6	Matrix 2	320.0	130.0	20.0	0	40	0.0	0	20	0.0	
Line Card #	1	2x10Gb/s	20.0	88.0	6.3	1	20	97.3	1	20	97.3	
Interface #	1	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	1	10	1.5	
Interface #	2	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	1	10	1.5	
Line Card #	2	2x10Gb/s	20.0	88.0	6.3	1	20	97.3	0	0	0.0	
Interface #	1	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	0	0	0.0	
Interface #	2	SFP 1x10G	10.0	1.4	0.1	1	10	1.5	0	0	0.0	

Table III. Breakdown of energy consumed by the emulated node ALUA in both the day and night configurations.

						DAY			Night			Saving
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #A		TSS320				1		**577.6**	1		**570.3**	**1.3%**
Common Parts							20	379.6		10	378.8	
Common Part #	1	FAN	320.0	30.0	0.0	1	20	30.0	1	10	30.0	
Common Part #	2	Supply	320.0	60.0	0.0	1	20	60.0	1	10	60.0	
Common Part #	3	EC	320.0	20.0	0.0	1	20	20.0	1	10	20.0	
Common Part #	4	SC	320.0	18.0	0.0	1	20	18.0	1	10	18.0	
Common Part #	5	Matrix 1	320.0	250.0	25.0	1	20	251.6	1	10	250.8	
Common Part #	6	Matrix 2	320.0	250.0	25.0	0	20	0.0	0	10	0.0	
Line Card #	1	10x10Gb/s	100.0	185.0	15.0	1	20	198.1	1	10	191.5	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Line Card #	2	10x10Gb/s	100.0	185.0	15.0	1	20	198.1	1	10	191.5	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	

Table IV. Breakdown of energy consumed by the emulated node ALUB in both the day and night configurations.

							DAY			Night		delta
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #B		TSS320				1		**577.6**	1		**0.0**	**100.0%**
Common Parts							20	379.6	0	0	0.0	
Common Part #	1	FAN	320.0	30.0	0.0	1	20	30.0	0	0	0.0	
Common Part #	2	Supply	320.0	60.0	0.0	1	20	60.0	0	0	0.0	
Common Part #	3	EC	320.0	20.0	0.0	1	20	20.0	0	0	0.0	
Common Part #	4	SC	320.0	18.0	0.0	1	20	18.0	0	0	0.0	
Common Part #	5	Matrix 1	320.0	250.0	25.0	1	20	251.6	0	0	0.0	
Common Part #	6	Matrix 2	320.0	250.0	25.0	0	20	0.0	0	0	0.0	
Line Card #	1	10x10Gb/s	100.0	185.0	15.0	1	20	198.1	0	0	0.0	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Line Card #	2	10x10Gb/s	100.0	185.0	15.0	1	20	198.1	0	0	0.0	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	

Table V. Breakdown of energy consumed by the emulated node ALUC in both the day and night configurations.

						DAY			Night			delta
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #C		TSS320				1		**577.6**	1		**570.3**	**1.3%**
Common Parts							20	379.6		10	378.8	
Common Part #	1	FAN	320.0	30.0	0.0	1	20	30.0	1	10	30.0	
Common Part #	2	Supply	320.0	60.0	0.0	1	20	60.0	1	10	60.0	
Common Part #	3	EC	320.0	20.0	0.0	1	20	20.0	1	10	20.0	
Common Part #	4	SC	320.0	18.0	0.0	1	20	18.0	1	10	18.0	
Common Part #	5	Matrix 1	320.0	250.0	25.0	1	20	251.6	1	10	250.8	
Common Part #	6	Matrix 2	320.0	250.0	25.0	0	20	0.0	0	10	0.0	
Line Card #	1	10x10Gb/s	100.0	185.0	15.0	1	20	198.1	1	10	191.5	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Line Card #	2	10x10Gb/s	100.0	185.0	15.0	1	20	198.1	1	10	191.5	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	0	0	0.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	0	0	0.0	

Table VI. Breakdown of energy consumed by the emulated node ALUD in both the day and night configurations.

						DAY			Night			Saving
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #D		TSS320				1		**879.8**	1		**894.4**	**-1.7%**
Common Parts							180	392.1		200	393.6	
Common Part #	1	FAN	320.0	30.0	0.0	1	180	30.0	1	200	30.0	
Common Part #	2	Supply	320.0	60.0	0.0	1	180	60.0	1	200	60.0	
Common Part #	3	EC	320.0	20.0	0.0	1	180	20.0	1	200	20.0	
Common Part #	4	SC	320.0	18.0	0.0	1	180	18.0	1	200	18.0	
Common Part #	5	Matrix 1	320.0	250.0	25.0	1	180	264.1	1	200	265.6	
Common Part #	6	Matrix 2	320.0	250.0	25.0	0	180	0.0	0	200	0.0	
Line Card #	1	10x10Gb/s	100.0	185.0	15.0	1	90	243.9	1	100	250.4	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Line Card #	2	10x10Gb/s	100.0	185.0	15.0	1	90	243.9	1	100	250.4	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	

Table VII. Breakdown of energy consumed by the emulated node ALUE in both the day and night configurations.

						DAY			Night			delta
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #E		TSS320				1		**879.8**	1		**894.4**	**-1.7%**
Common Parts							180	392.1		200	393.6	
Common Part #	1	FAN	320.0	30.0	0.0	1	180	30.0	1	200	30.0	
Common Part #	2	Supply	320.0	60.0	0.0	1	180	60.0	1	200	60.0	
Common Part #	3	EC	320.0	20.0	0.0	1	180	20.0	1	200	20.0	
Common Part #	4	SC	320.0	18.0	0.0	1	180	18.0	1	200	18.0	
Common Part #	5	Matrix 1	320.0	250.0	25.0	1	180	264.1	1	200	265.6	
Common Part #	6	Matrix 2	320.0	250.0	25.0	0	180	0.0	0	200	0.0	
Line Card #	1	10x10Gb/s	100.0	185.0	15.0	1	90	243.9	1	100	250.4	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Line Card #	2	10x10Gb/s	100.0	185.0	15.0	1	90	243.9	1	100	250.4	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	

Table VIII. Breakdown of energy consumed by the emulated node ALUF in both the day and night configurations.

						DAY			Night			delta
	#	Desc.	Par. #1 (Rate Gb/s)	Power (static)	Power (dinamic)	ON / OFF	Current Rate (Gb/s)	Power	ON / OFF	Current Rate (Gb/s)	Power	
NODE #F		TSS320				1		**879.8**	1		**894.4**	**-1.7%**
Common Parts							180	392.1		200	393.6	
Common Part #	1	FAN	320.0	30.0	0.0	1	180	30.0	1	200	30.0	
Common Part #	2	Supply	320.0	60.0	0.0	1	180	60.0	1	200	60.0	
Common Part #	3	EC	320.0	20.0	0.0	1	180	20.0	1	200	20.0	
Common Part #	4	SC	320.0	18.0	0.0	1	180	18.0	1	200	18.0	
Common Part #	5	Matrix 1	320.0	250.0	25.0	1	180	264.1	1	200	265.6	
Common Part #	6	Matrix 2	320.0	250.0	25.0	0	180	0.0	0	200	0.0	
Line Card #	1	10x10Gb/s	100.0	185.0	15.0	1	90	243.9	1	100	250.4	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Line Card #	2	10x10Gb/s	100.0	185.0	15.0	1	90	243.9	1	100	250.4	
Interface #	1	SFP 1x10G	10.0	4.4	0.6	0	0	0.0	1	10	5.0	
Interface #	2	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	3	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	4	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	5	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	6	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	7	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	8	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	9	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	
Interface #	10	SFP 1x10G	10.0	4.4	0.6	1	10	5.0	1	10	5.0	

3.2.2 Router Network Cloud

This section reports the results of the router cloud with a particular focus on the internal behavior of the DROP prototypes, and on the two emulation networks running the ESIR and the GRiDA routing mechanisms, respectively.

It is worth recalling that section 6.2 of D6.4 reports all the details of the internal architecture, the building blocks, the green mechanisms integrated and the performance indexes exposed by all the physical and virtual nodes of the core cloud. In order to correctly understand the results and the discussions reported in the following, it is warmly suggested to read section 6.2 in D6.4 before proceeding to the next subsections.

3.2.2.1 Physical nodes CNIT1 and CNIT2

The following internal measurements are reported for the CNIT1 and CNIT2 routers:

- Energy consumption of the routers according to the time slot broken down on a per DROP element basis (CNIT1 in Fig. 45, CNIT2 in Fig. 49);
- Traffic volumes sent to the Forwarding Elements (FEs) according to the time slot (CNIT1 in Fig. 46, CNIT2 in Fig. 50);
- Number of cores left active in the server-based FEs according to the time slot (CNIT1 in Fig. 47, CNIT2 in Fig. 51);
- Operating frequency values and supply voltage of the cores of the Broadcom XLP832 System on Chip (SoC) in the FE3 according to the time slot (CNIT1 in Fig. 48, CNIT2 in Fig. 52).

Concerning the energy consumption depicted in Fig. 45 and Fig. 49, we can clearly note how the power profile of both DROP instances follows the traffic volume trend, by providing ~52% of saving. This proportionality between energy consumption and traffic load has been obtained thanks to the device-level and element-level LCPs (see D6.2 and D6.4 for further details).

In particular, the device-level LCP is responsible for moving traffic flows among the FEs and putting the unused ones into sleep mode, while the element-level LCP controls the power scaling configuration of each forwarding element. Fig. 46 and Fig. 50 show the traffic share among FEs according to the time slots for the CNIT1 and the CNIT2 routers, respectively. As one can observe, after ~15 time slots the device-level LCP of both routers decides to put FE2 into sleep mode, and accordingly the traffic is migrated among the remaining active FEs.

During the low traffic period (between the 50th and the 200th time slots), also FE1 is bypassed and the global power consumption is dominated by the XLP and the interconnection elements only. However, Fig. 49 highlights a non-negligible power contribution of FE2 also when it is in stand-by mode. This effect is because FE2 of the CNIT2 router has a relatively high power dissipation (around 10 Watt) also in sleep mode.

Then, after the “night hours”, the device-level LCP starts waking up the FEs according to the increase of the offered load and the consumption reaches a peak around 500 W.

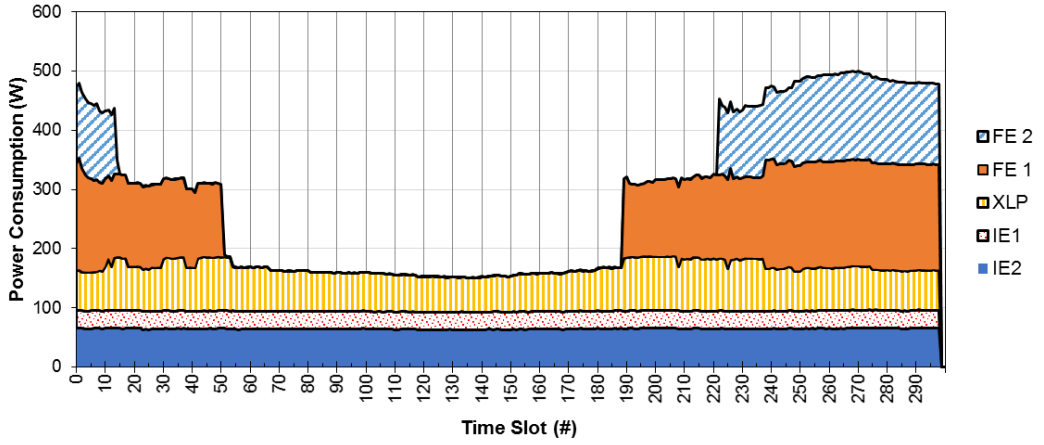

Fig. 45. Energy consumption of the CNIT1 router according to time slots. The results are visualized in terms of stacked areas in order to express the total energy consumption of the prototype.

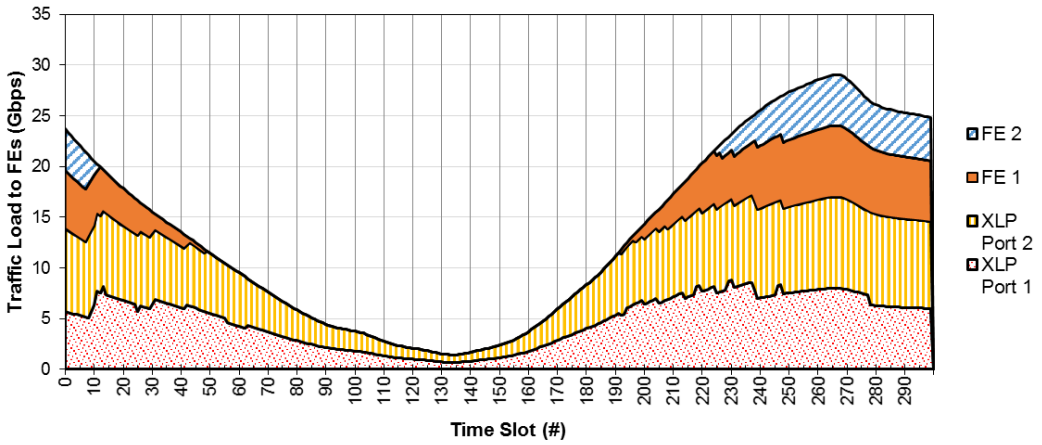

Fig. 46. Volume of traffic sent to the FEs of the CNIT1 router according to time slots. The results are visualized in terms of stacked areas in order to express the total offered traffic load to the prototype.

Regarding the element-level LCPs, two different optimization strategies have been developed in the context of WP 3 and WP 6: one for the FEs based on commercial off-the-shelf servers and one for the evaluation board of the Broadcom XLP832.

The former selects the number of cores needed for forwarding and puts the remaining ones in a deep sleeping state depending on the current incoming load [3]. The latter exploits the voltage and frequency scaling techniques for adapting the processor capacity to the current incoming traffic. Fig. 47 and Fig. 51 show the number of active cores of the FEs for the two DROP routers, respectively. The number of cores is 0 when the forwarding elements enter sleep mode during the low traffic period.

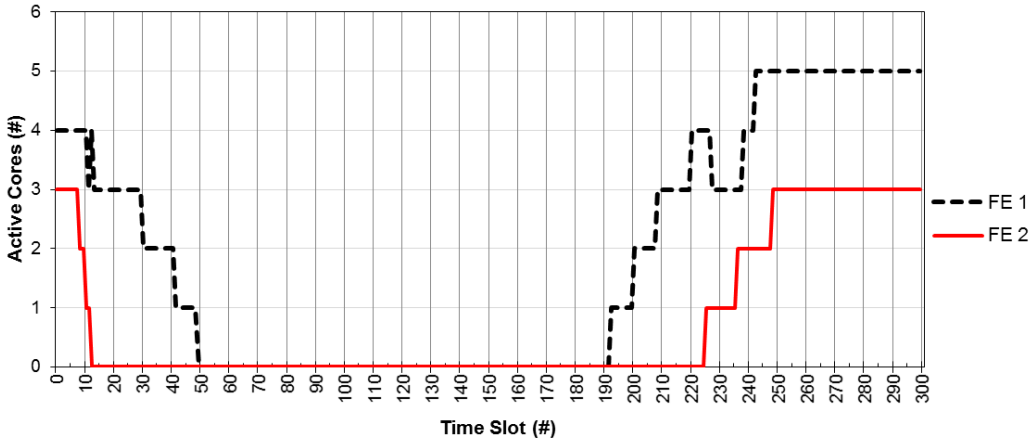

Fig. 47. Number of active cores in FE1 and FE2 of the CNIT1 router according to time slots. When no cores are activated, the entire FE is put into standby mode.

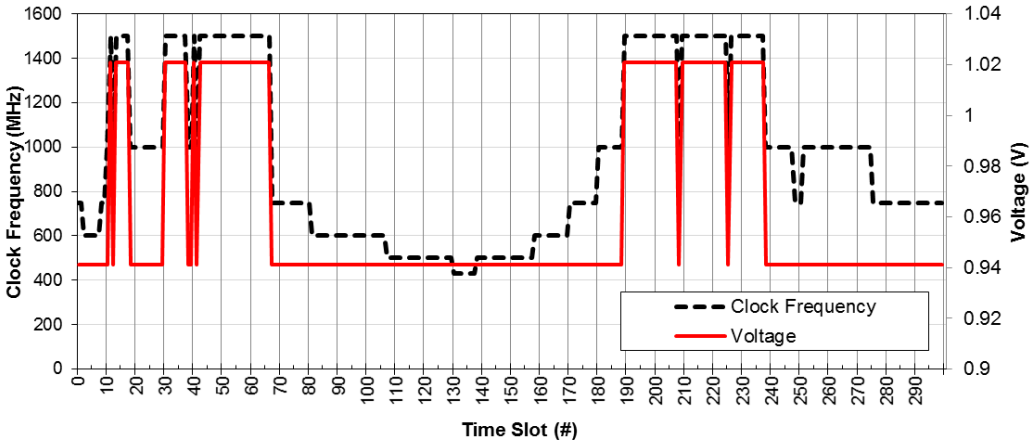

Fig. 48. Frequency and supply voltages of the cores in the FE3 (based on the Broadcom XLP SoC) of the CNIT1 router according to time slots. The frequency and the voltage values are kept uniform in all the cores.

Regarding the XLP, Fig. 48 and Fig. 52 depict the working clock frequency and supply voltage according to the time slots. The maximum supply voltage is selected when the LCP sets the maximum working frequency (see Fig. 48). In the CNIT2 router case (Fig. 52), the supply voltage remains constant and only the clock frequency changes over the time. This is because in the CNIT2 router the XLP receives less traffic and the maximum working frequency is never chosen.

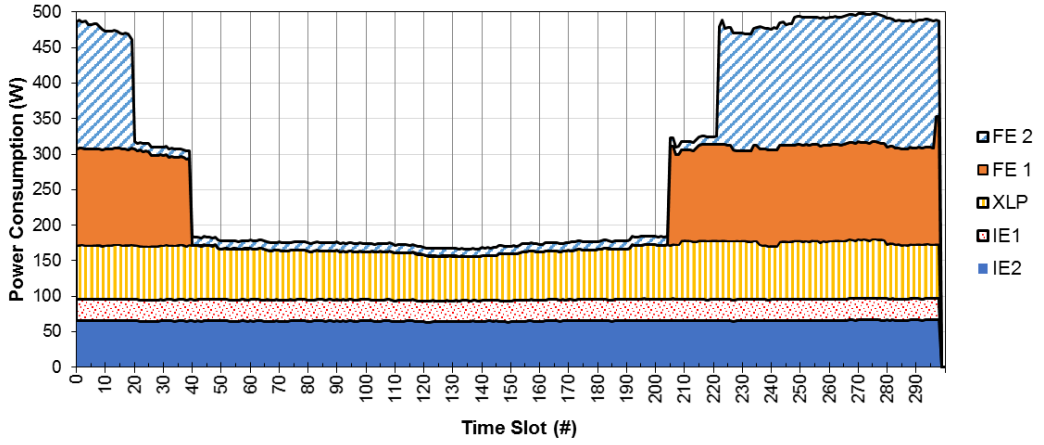

Fig. 49. Energy consumption of the CNIT2 router according to time slots. The results are visualized in terms of stacked areas in order to express the total energy consumption of the prototype.

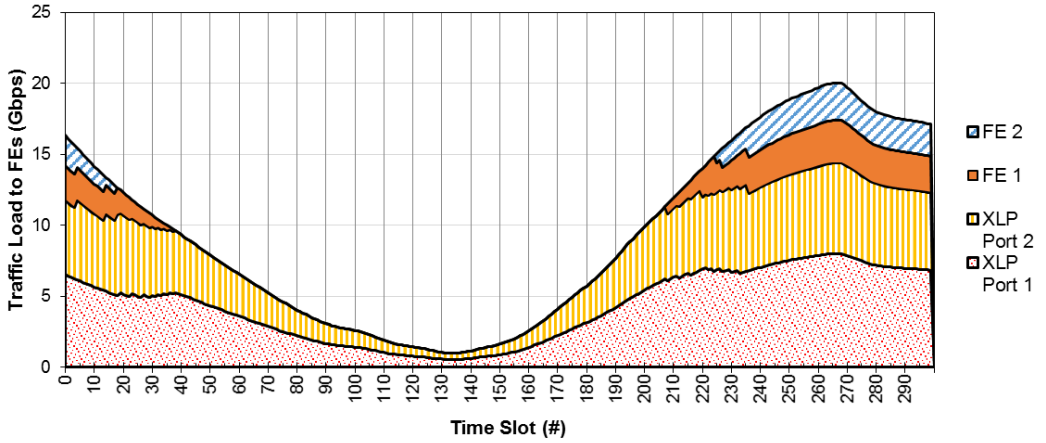

Fig. 50. Volume of traffic sent to the FEs of the CNIT2 router according to time slots. The results are visualized in terms of stacked areas in order to express the total offered traffic load to the prototype.

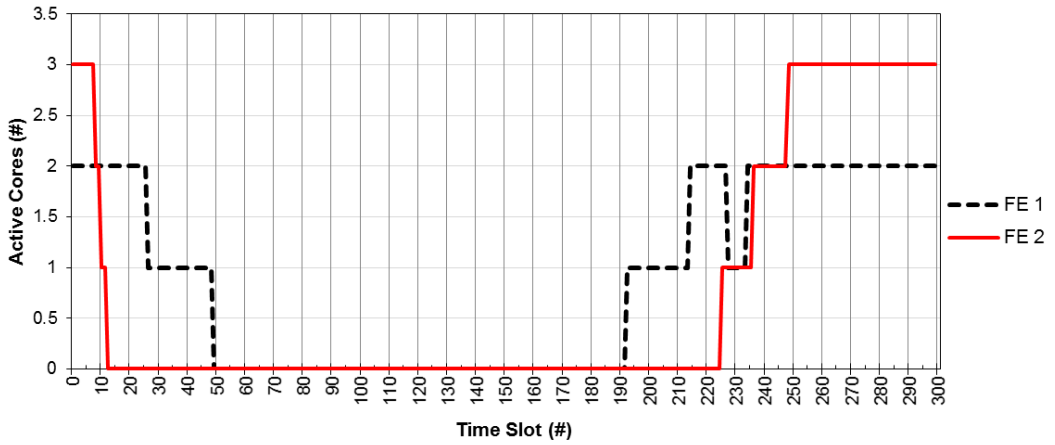

Fig. 51. Number of active cores in FE1 and FE2 of the CNIT2 router according to time slots. When no cores are activated, the entire FE is put into standby mode.

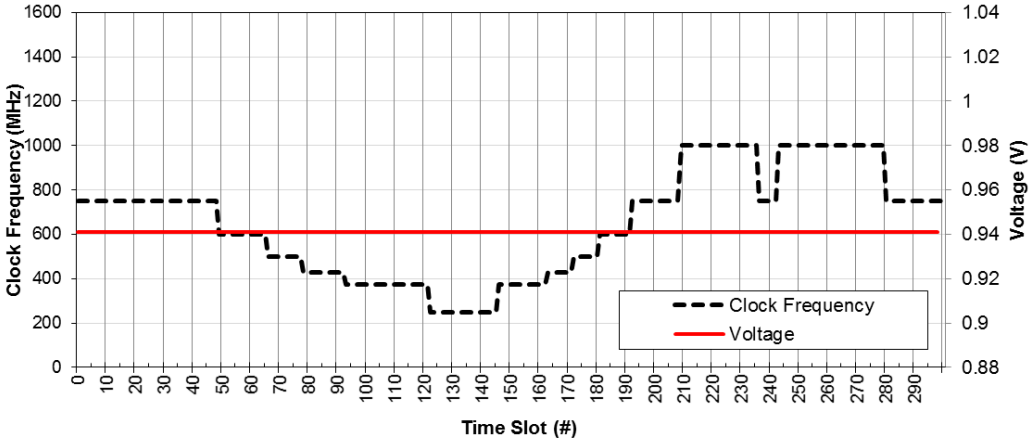

Fig. 52. Frequency and supply voltages of the cores in FE3 (based on the Broadcom XLP SoC) of the CNIT2 router according to time slots. The frequency and the voltage values are kept uniform in all the cores.

The above experimental results demonstrate that the power management capabilities developed in the DROP router prototype guarantee a good level of energy proportionality with respect to the traffic load. The results are even more convincing considering that no service degradation (e.g., delays or losses) has been experienced during the demonstration.

3.2.2.2 Emulation Network 1

The emulation network 1 in the router cloud was introduced and completely defined in Sections 2, 3 and 6.2.2 in the D6.4 report. It runs the ESIR routing algorithm over a network of 38 nodes composed of 36 virtual machines, of the CNIT1 and CNIT2 routers and of 144 links (see Figure 4 in the D6.4 report). The CNIT1 and CNIT2 physical prototypes have been mapped onto the nodes #25 and #26 of the topology.

The traffic matrix is composed by 154 flows, whose offered load is fluctuating over the day according to Figure 12 in the D6.4 report.

Fig. 51 shows the percentage of links in standby mode for the different time slots. In particular two cases are considered:

- the bi-directional case, i.e. the standby of a network link is possible only if both unidirectional links composing the bi-directional link are put in standby
- the unidirectional case, i.e. the standby can be enabled independently on each unidirectional link

Fig. 51 shows that the percentage of links put in standby varies from about 15% and 6% during low traffic hours, to 8% and 2%, during peak traffic hours, for the unidirectional and bi-directional cases, respectively. The ability of ESIR to adapt the network configuration to traffic levels is clear. The results also highlight that the performance of ESIR in the unidirectional case is substantially better than the one of the same algorithm in the bi-directional case: in almost all time slots the percentage of standby links in the unidirectional case is more than twice the value obtained in the bi-directional case. The main reason behind this behaviour is that the unidirectional case was the scenario considered in the ESIR routing algorithm when it was defined: the adoption in the bi-directional case was obtained by executing ESIR in the unidirectional case and then introducing the bi-directional

constraint on link standby. To improve the performance of ESIR in the bi-directional case, a deep modification of the algorithm would be required: in particular, a new definition of the "move" concept, that represents the building block of ESIR, should be provided.

Figs. 52 and 53 report the maximum and average utilization of network links over the day for bi-directional and unidirectional cases. The results concerning the average link utilizations are quite similar in both cases: lower values are obtained for the bi-directional case since ESIR is able to put a higher number of links in the unidirectional case to sleep, as explained before. This behavior is much more evident by evaluating the maximum link utilization: the values obtained for the unidirectional cases are substantially higher than the ones of the bi-directional case. Anyway, Figs. 52 and 53 show one of the main features of ESIR, the ability to maintain the maximum link utilization under a fixed threshold value: in this case 50% has been fixed as threshold value.

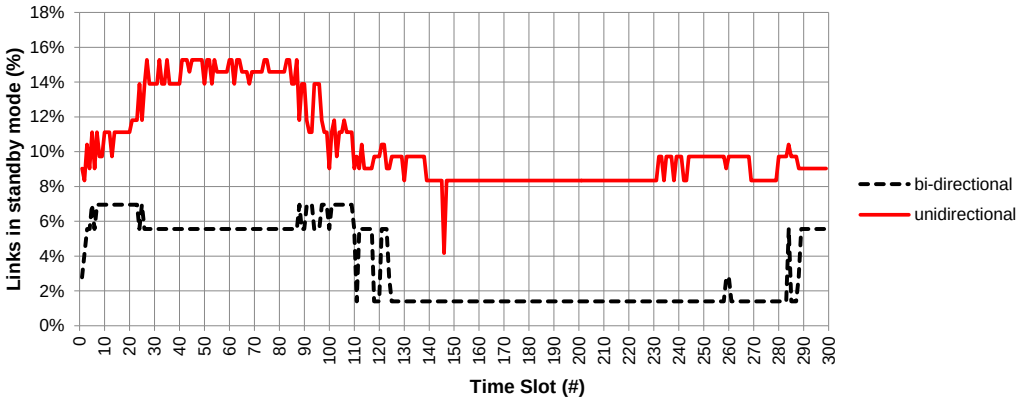

Fig. 53. Percentage of the links put into standby modes according to time slots in the case with unidirectional links and the one with bi-directional links.

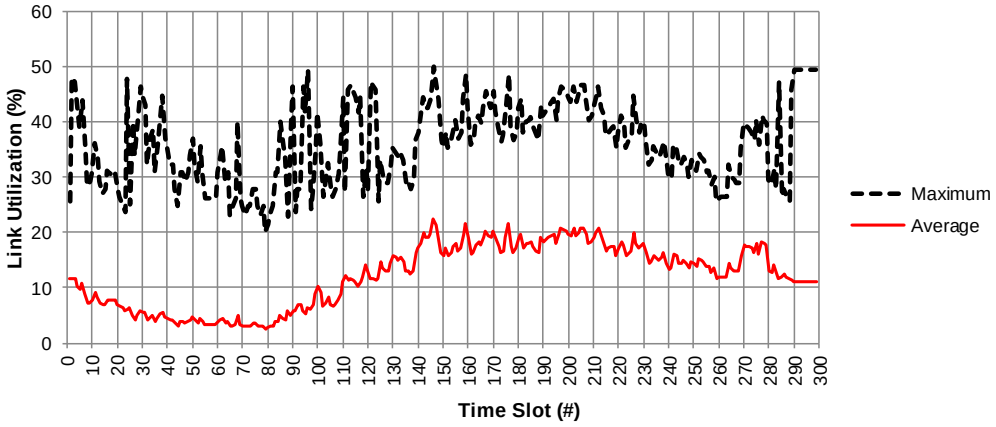

Fig. 54. Maximum utilization of network links in the unidirectional case according to time slots.

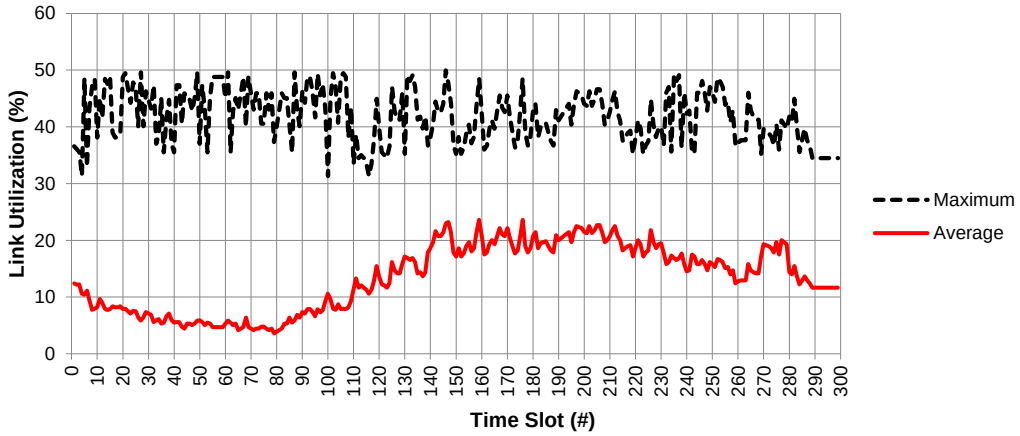

Fig. 55. Maximum utilization of network links in the bidirectional case according to time slots.

3.2.2.3 Emulation Network 2

The emulation network 2 in the router cloud was introduced in Sections 2, 3 and 6.2.3 of the D6.4 report. It runs the GRiDA routing algorithm over a network of 6 nodes composed of 4 virtual machines, of the CNIT1 and CNIT2 routers and of 7 links (see Figure 5 in the D6.4 report).

Fig. 56, which is also present in the D6.4 report, shows the offered load of the flow crossing the network between nodes 1 and 5. This figure has been re-introduced here for easing the understanding of the obtained results.

Figs. 57-63 report the bandwidth utilization according to time slots for all the emulated network links. As can be noted, routing re-configurations happen close to the time instant in which the traffic source changes the traffic load level (i.e., slots #40, #160, #220, #260).

As shown in Fig. 64 and as evident from the link utilizations in Figs. 57-63, the network routing converges to three different configurations according to the offered load level. In more detail, when the traffic load is low (4 Mbps between slots #40 and #160), the path of the emulated flow only includes the EM2-L-3 and EM2-L-6 links.

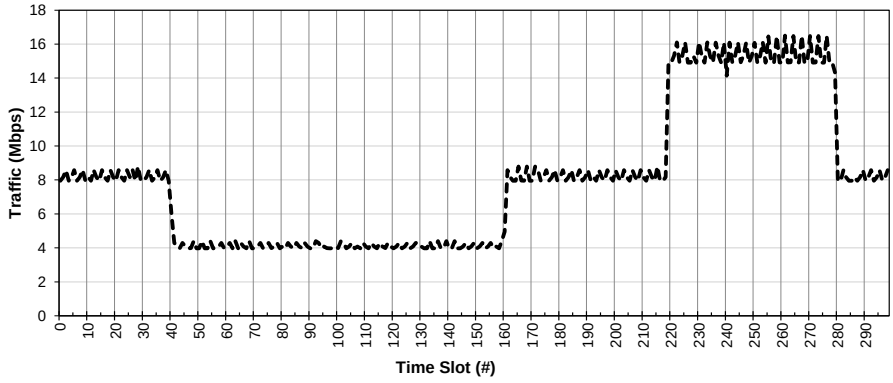

Fig. 56. Offered load of the emulated traffic flow crossing the emulation network 2 according to time slots.

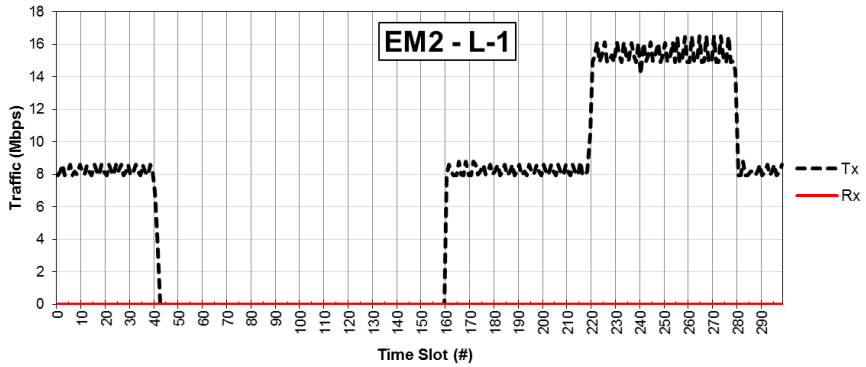

Fig. 57. Bandwidth occupancy according to time slots of link L-1 of the Emulation Network 2 in the Core Cloud (maximum bandwidth 30 Mbps).

When the traffic level rises to 8 Mbps (between slots #0 and 40, #60 and #220, and #280 and #300), given the limited bandwidth of the EM2-L-3 link, the path is changed to the EM2-L-1, EM2-L-4 and EM2-L-6 links. As shown in Fig. 65, the energy consumption of the network increases, but the new path allows avoiding traffic congestion and packet losses.

Finally, when the traffic level achieves 16 Mbps, the EM2-L-4 link becomes a bottleneck. So, the GRiDA algorithm decides to further update the routing configuration, passing to a more energy-consuming one. The flow path in the new configuration passes through the EM2-L-1, EM2-L-2, EM2-L-5 and EM2-L-6 links.

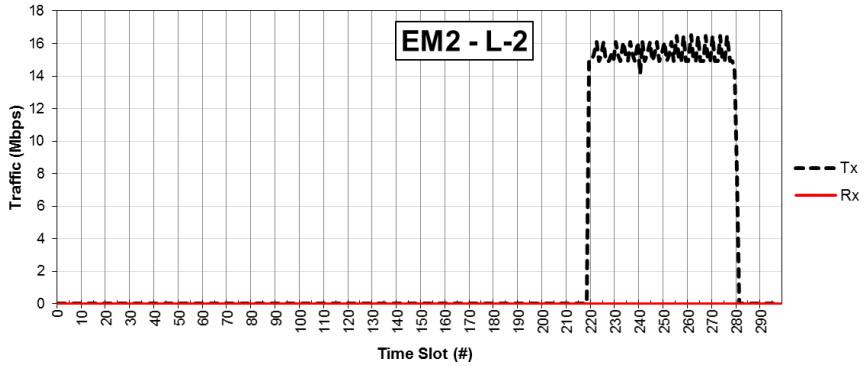

Fig. 58. Bandwidth occupancy according to time slots of link L-2 of the Emulation Network 2 in the Core Cloud (maximum bandwidth 30 Mbps).

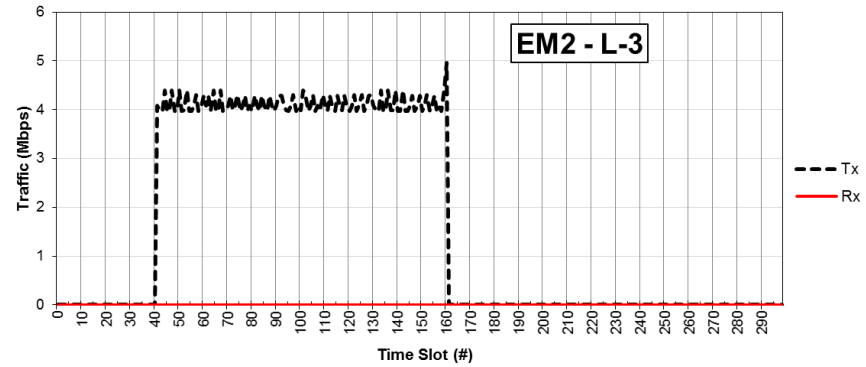

Fig. 59. Bandwidth occupancy according to time slots of link L-3 of the Emulation Network 2 in the Core Cloud (maximum bandwidth 5 Mbps).

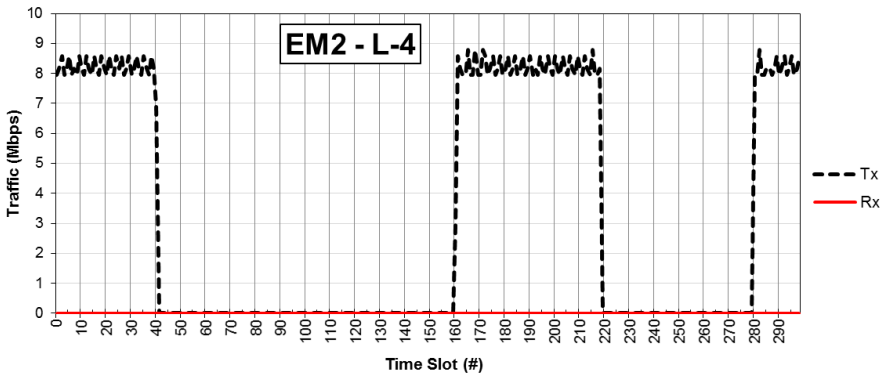

Fig. 60. Bandwidth occupancy according to time slots of link L-4 of the Emulation Network 2 in the Core Cloud (maximum bandwidth 10 Mbps).

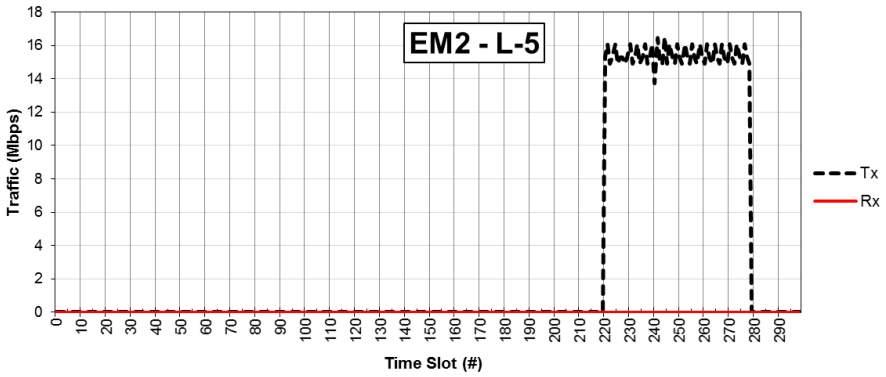

Fig. 61. Bandwidth occupancy according to time slots of link L-5 of the Emulation Network 2 in the Core Cloud (maximum bandwidth 30 Mbps).

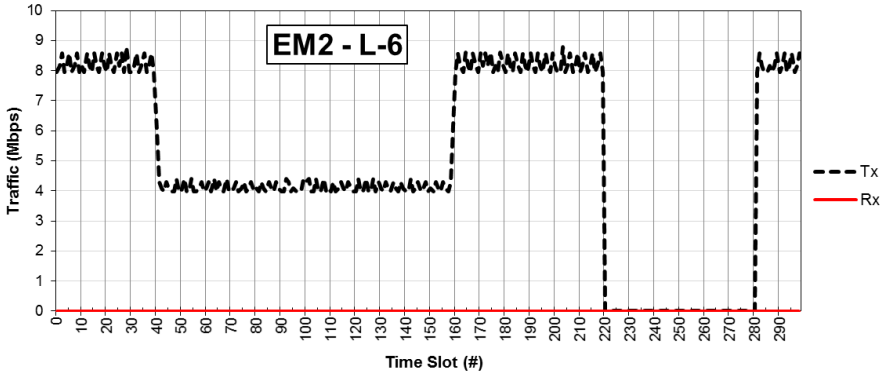

Fig. 62. Bandwidth occupancy according to time slots of link L-6 of the Emulation Network 2 in the Core Cloud (maximum bandwidth 10 Mbps).

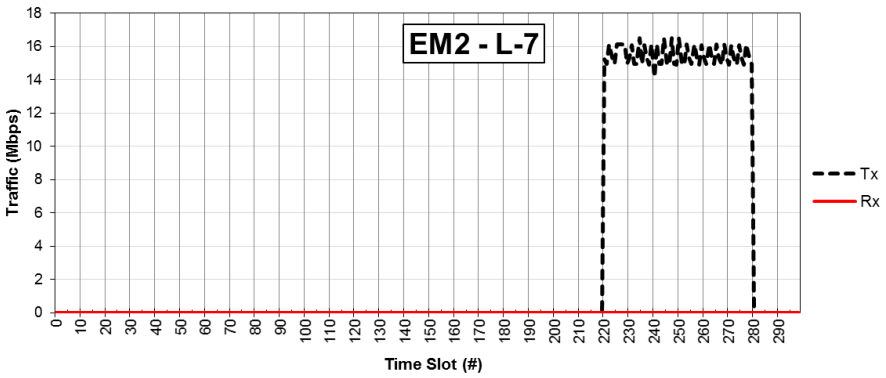

Fig. 63. Bandwidth occupancy according to time slots of link L-7 of the Emulation Network 2 in the Core Cloud (maximum bandwidth 30 Mbps).

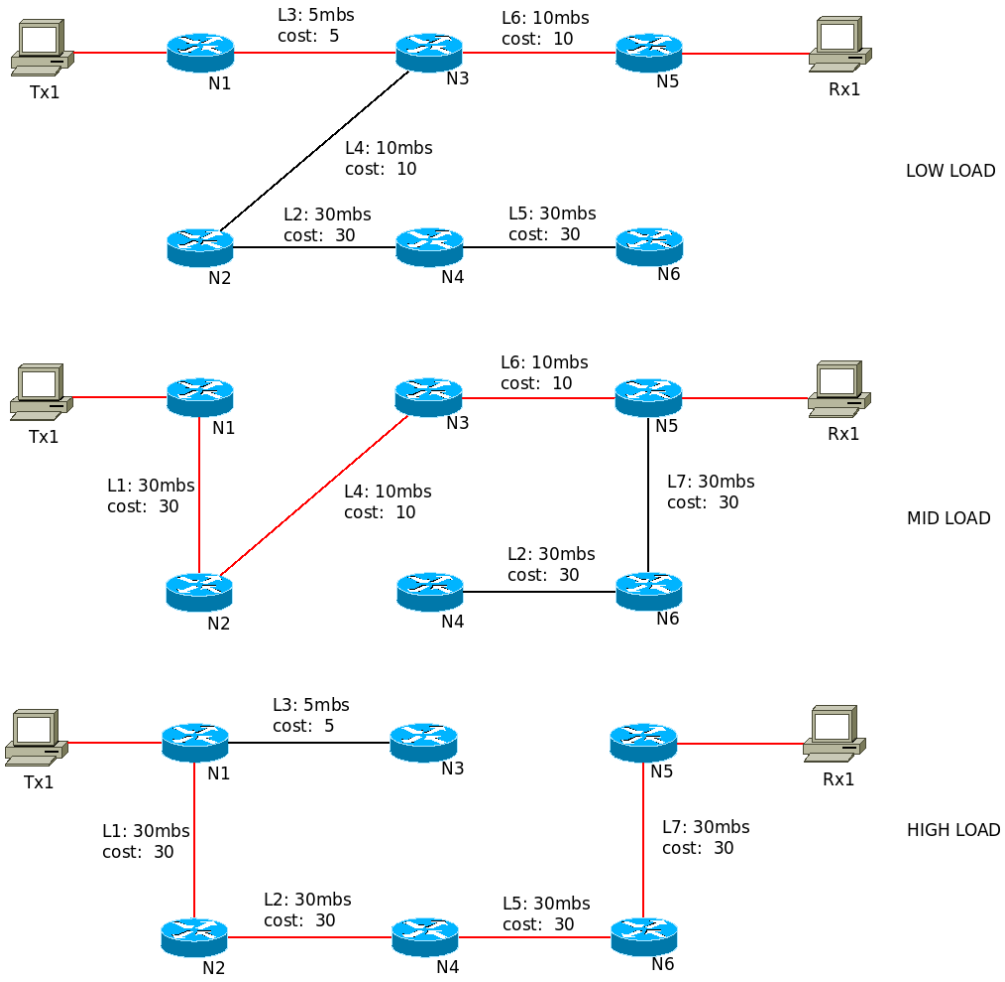

Fig. 64. Routing configurations according to the traffic offered load to which the GRiDA algorithm converges.

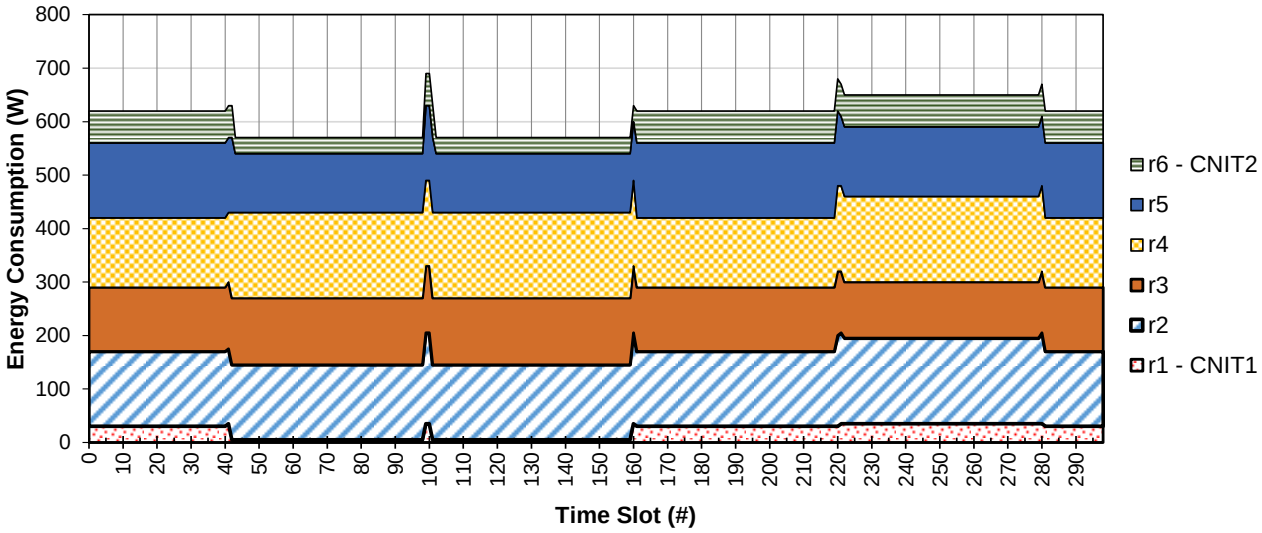

Fig. 65. Energy consumption of the Emulation Network 2. For the physical CNIT1 and CNIT2 routers (mapped on nodes 1 and 2), only the "emulated" energy consumption of the interfaces towards the emulation network 2 has been reported.

3.2.3 Metro Network Cloud

The test has been executed according to test setup and Energy Aware features described in D6.4, Section 5.3 “Metro Transport”. A Metro Ring was emulated by means of 3 devices (Node A, Node B, Node C with reference to Fig. 66, or TEI1, TEI2 and TEI3 with reference to Fig. 2 and 3), where Node A and Node B support Energy Aware features, while Node C is a Business-as-Usual platform.

Although slightly suboptimal in terms of maximum savings that can be achieved, having Node C as Business-as-Usual platform demonstrates that the proposed mechanisms and solutions can work well also in the presence of heterogeneous energy capabilities and profiles at network nodes, and that it is not fundamental to update the entire network infrastructure to achieve significant results. Indeed, Operators can start saving a significant amount of energy (and, then, of operating expenses) from the first Energy Aware equipment installed.

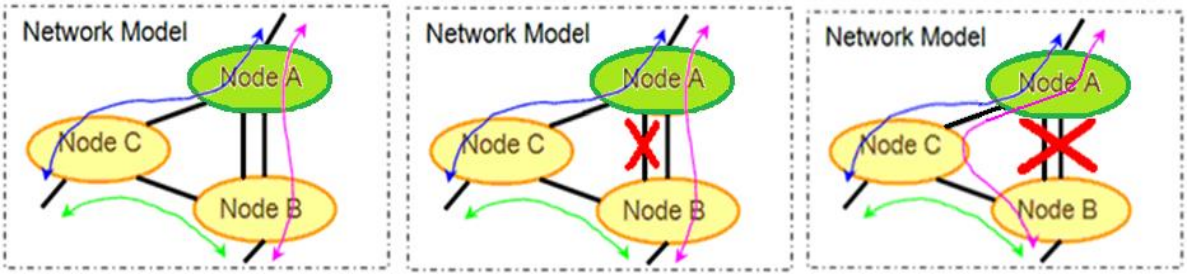

Fig. 66. The reference model of the network-wide strategy adopted in the metro network ring.

Network-wide control strategies (i.e., routing and traffic engineering) give the possibility of moving traffic load among network nodes. When a network is under-utilized, network load can be moved to few “active” nodes, and all the others can enter standby modes.

In our demo, the IP Layer is overlaid over an optical network. At layer 2, Ethernet frames are used to map IP traffic on the physical infrastructure. Each physical (optical) channel can carry multiple “virtual” L2 links (or Virtual LAN, VLANs according to IEEE 802.1q), which directly connect two nodes working at L3 level.

In this way, we fully exploit the L2 protocol switching features to migrate L2 VLANs from one link or card going into sleeping (or power-off) mode to an alternative path remaining active. It is worth noting that, thanks to this approach, any modifications to the IP logical layer (i.e., to the IP topology) are avoided, and no re-convergence of L3 routing protocols are actually needed [4]. In other words, we apply re-mapping possibilities for a set of VLAN, while the IP mapping is kept fixed, in the sense that the neighbouring L3 nodes always see the same L3 interfaces and logical links.

It can be noted that this approach is indeed the same that is applied in case of fault events, except for the fact that, in our case, the L2 re-mapping applies before the link or the card becomes unavailable. Therefore, by means of a careful synchronization of traffic routing and sleeping/wake-up events, energy saving mechanisms can be applied without any traffic losses or service interruptions.

This synchronization is made easier by the presence of a centralized Network Control Process (NCP) that is constantly monitoring the network and/or node configurations and the traffic loads. The NCP has the knowledge of the network topology, and applies topology modifications (like, for instance, waking up or putting network resources to sleep), as soon as traffic loads exceed or fall

below pre-defined thresholds. In order to avoid network topology fluctuations or instabilities, a suitable hysteresis mechanism is applied at each traffic threshold.

A further development activity (left to future studies), which might significantly increase the efficiency of the mechanisms here proposed, regards the dynamic adaptation of these thresholds.

During the demonstration, traffic load varies according to the given profile (see Sect. 3 of the D6.4 report), and Network topology changes step by step according to Fig. 67 (the behaviors of the Energy Aware Link Aggregation Groups (LAGs) are not shown for the sake of simplicity).

Fig. 68 shows the number of active links on L-12 and L-13 Energy-Aware LAGs according to the traffic load. It can be highlighted how the number of active links follows the carried traffic closely. Moreover, when VLAN remapping happens to the F-6 flow and the L-12 Energy-Aware LAG is entirely put into standby mode (at time slot #80), it can also be seen how the traffic load climbs up on the L-13 LAG (since the F-6 flow is redirected to this link). Similarly, at time slot #180, when the L-12 link is re-activated, the traffic load on the L-13 LAG falls suddenly down, since F-6 is migrated back to the L-12 link.

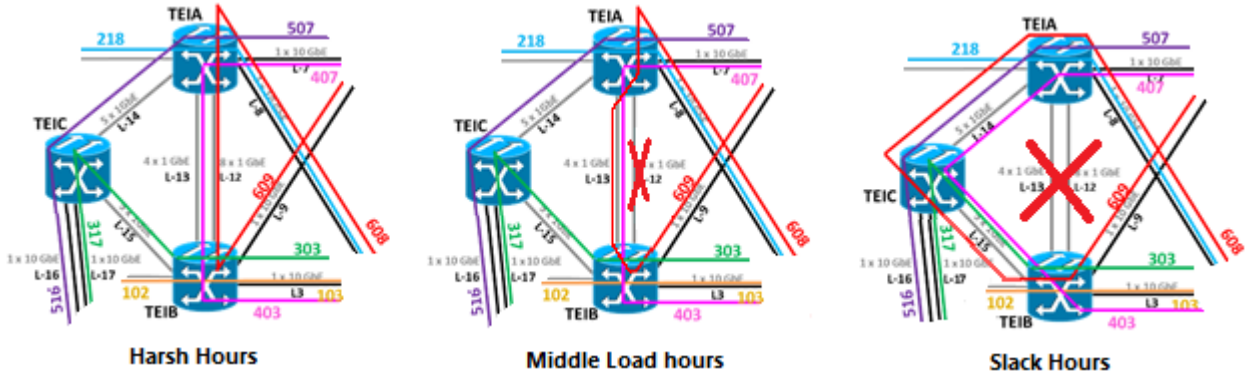

Fig. 67. The three metro cloud configurations used during harsh hours, middle load hours, and slack hours.

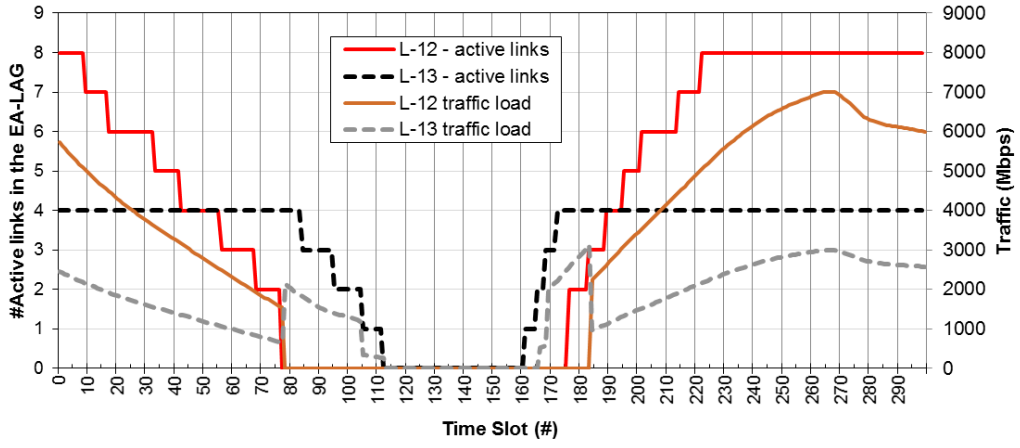

Fig. 68. Number of active links in the energy-aware LAGs (L-12 and L-13 with respect to the overall ECONET demo topology) and volume of carried traffic on those links.

Fig. 69 shows the overall power consumption measured at each time slot, and it makes evident that the Business-As-Usual device (i.e., Node C) exhibits an almost constant energy requirement, since there are no visible variations according to the fluctuating traffic load.

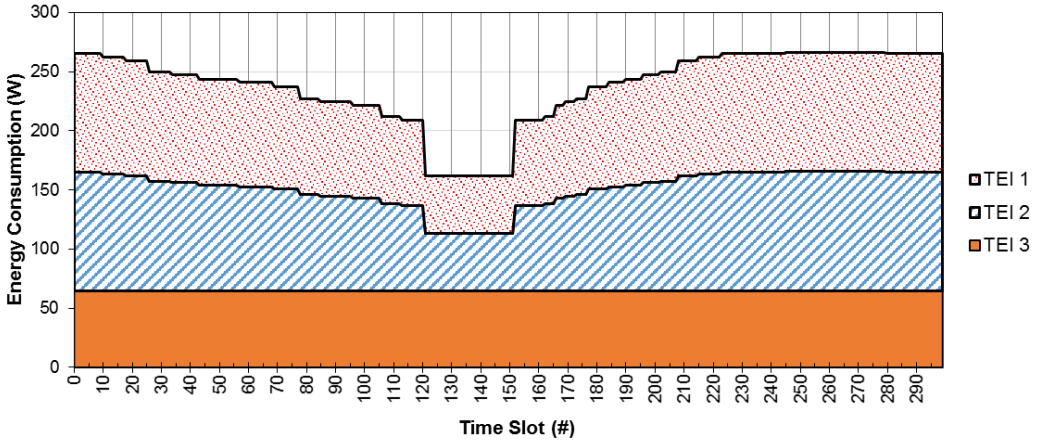

Fig. 69. Overall energy consumption of the metro cloud according to time slots. The energy consumption is broken down per device.

According to Fig. 69 and as expected, most of the power reduction of the energy-aware devices is concentrated during slack hours (from time slot #120 to time slot #150), where savings up to 52% of the Business-as-Usual energy requirements are achieved.

On average, the measured level of power consumption savings is ~14.3%. This is a very good and encouraging result, although greatly suboptimal. In fact, in our proof-of-concept demonstration, typical network dimensioning criteria were not considered, such as:

- Peak traffic load estimation,
- Over-provisioning from 25% to 50% in respect to estimated peak load,
- Redundancy for fault recovery.

As an example, assuming a peak rate of ~7 Gbps, an over-provisioning calculated over 10 years of lifetime, and typical redundancy configurations, the typical provisioned link capacity would be 20 Gbps. Therefore, according to the day-night traffic variations, average traffic utilizations can be estimated to be approximately 3 Gbps, which corresponds to ~ 15% of available link capacities. By extending this reasoning to the entire network, it gives clear indication that ECONET and Energy Aware Networking hold the promise to enable extremely aggressive power saving, well above 50%.

During the tests, we did not experience any packet loss, as we did not experience any IP topology re-mapping. VLAN re-mapping times according to the number of L2 virtual links per physical link and to the number of flow table entries to be updated during the migration process have been measured to span a range from 20 to 200 ms.

3.2.4 Datacenter Network Cloud

The data center network emulation was done via the VPI/Ethernet switch SX1036. This switch incorporates 36 ports that were connected via snake topology to the Metro cloud on one side and to the emulated server on the other side, as described in document D6.4. This switch is capable to work in two power modes to account for different flavors of aggressiveness of power reduction features:

Power mode features incorporated into Power_mode_1 do not affect the link bandwidth, like smart clock gating, voltage/frequency scaling, smart thermal algorithm, etc.

In Power_mode_2 all the above power features are incorporated and the Width Reduction Power Scaling feature (WRPS) is operational as well.

The control of shifting the operation between Power_Mode_1 and Power_Mode_2 is done via GAL commands. Moreover, the switch incorporates power monitoring features to be able to track the system power consumption from NCP devices by the GAL REST protocol.

During the ECONET demo, all the above capabilities were demonstrated.

The Data Centre Network emulation power is depicted in the Fig. 70:

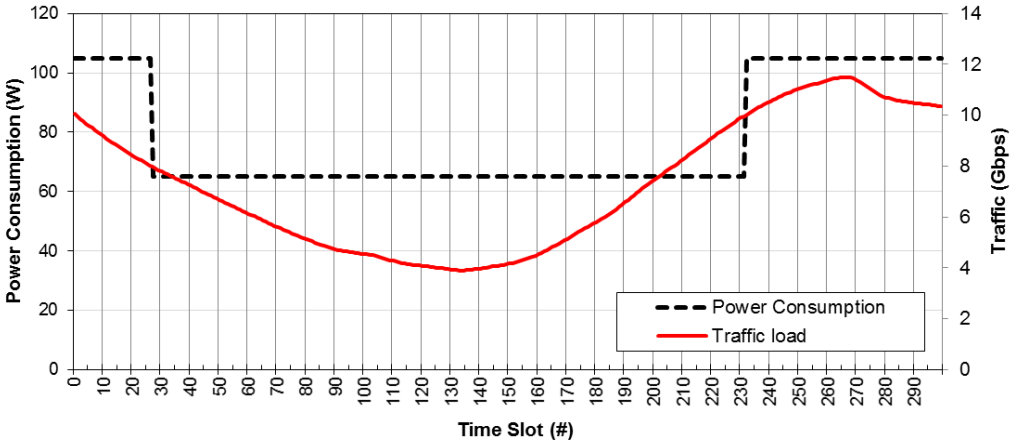

Fig. 70. Energy consumption of the MLX1 prototype and traffic load on the links belonging to the "snake" topology according to time slots.

As can be seen, the emulated network power is changing between ~105 W and ~65 W, resulting in 38% reduction in the network power. The high power is caused by tracking high bandwidth on the links of over 10 Gbps, up to 12 Gbps (from transitioning from 1X operation to 4X operation), while the low power is caused when moving to the lower link bandwidth of under 8 Gbps, down to 4 Gbps (due to transitioning from 4X operation to 1X operation). It should be noted that since the traffic pattern in the demo has not dropped all the way to 0 bps bandwidth, and especially since the ports have not been put into entire shut-down state, then the full power reduction capability of the switch was not captured. This capability was demonstrated on the stand-alone case and was described in document D3.3. An NCP can choose to move traffic from non-loaded ports and in that case to be able to shut down ports with 0 bps traffic, fully exploiting the shutdown power features incorporated into the switch.

Since in the ECONET demonstration the entire link bandwidth in the snake topology was throttling between 12 Gbps and 4 Gbps, and since the ports of the switch can deliver 40 Gbps, we conducted further tests to capture the entire range of power reduction on the above two modes.

Figure 71 captures the system power of the switch in Power_Mode_1 when throttling the bandwidth between the full capability ranges (10Gbps – 40Gbps):

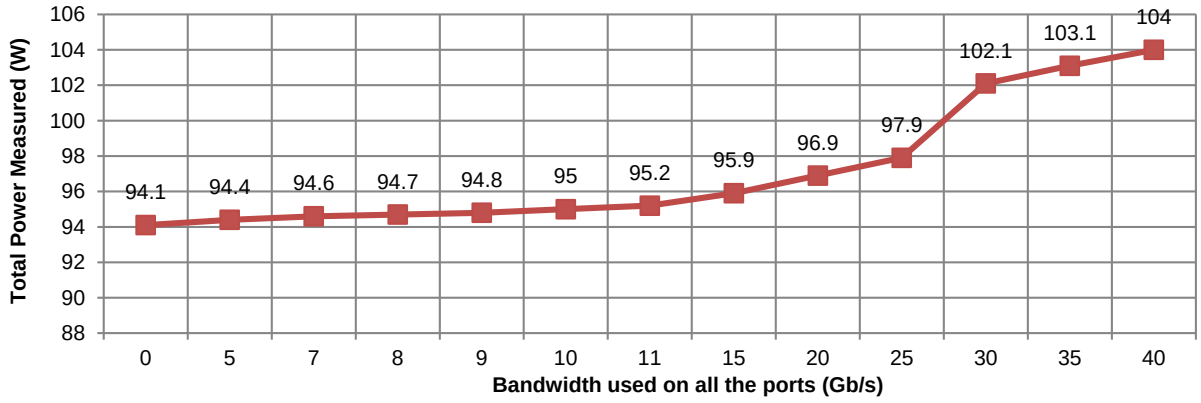

Fig. 71. System Power_mode_1: WRPS DIS.

The above power reductions are caused mainly due to smart clock gating, frequency/voltage scaling algorithm and Smart Thermal Optimization algorithm, all of which are described in D3.3. The transition between the modes and the power monitoring is performed by incorporation of the GAL REST API, as described in D5.5. Those measurements were taken in a thermal-controlled environment. Figure 72 shows the Fan speed change during the above test.

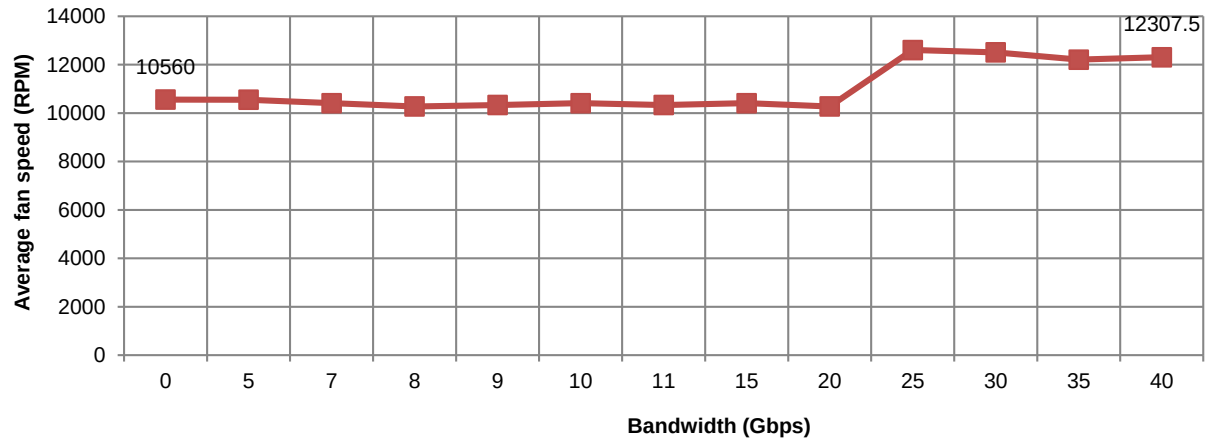

Fig. 72. Average Fan Speed.

As can be seen, the Fan speed is adjusted based on the used power of the system and it shifts between discrete ranges generated by closed loop-tracking the hysteresis band of the thermal conditions of the system. This demonstrates the Smart Algorithm Thermal tracking mechanism described in D3.3.

Similarly, Fig. 73 depicts the reduction in power when using the more aggressive power mode of the switch, i.e. Power_Mode_2:

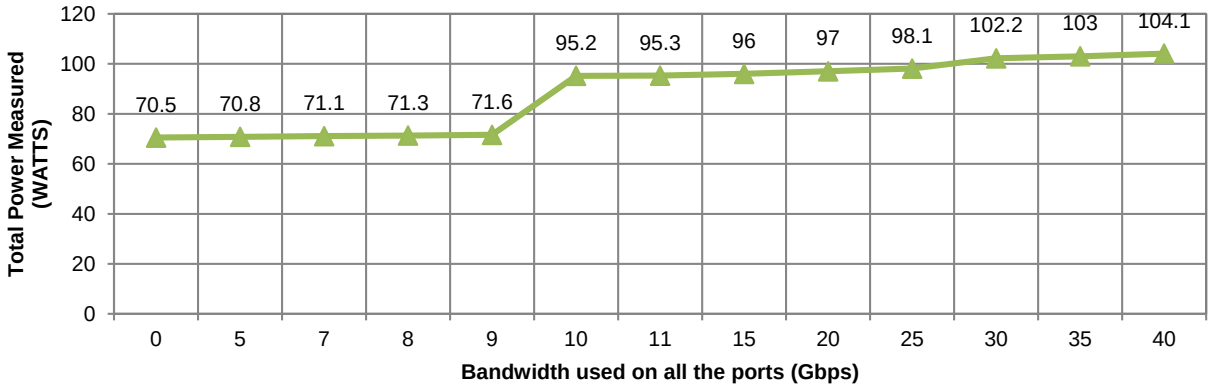

Fig. 73. Energy consumption when the WPRS feature is enabled.

The results in Fig. 73 demonstrate the adaptive algorithms implemented in Mellanox devices to optimize the power consumption based on activity and bandwidth used on the links.

3.2.5 Access Network Cloud

Fig. 74 reports the number of active ports left active and the energy consumption of the ETY1 and ETY2 prototypes according to time slots.

As can be noted and as expected, the shapes of these curves follow closely the ones of the offered load of the F-1 and F-2 flows, which enter ETY2 and ETY1, respectively.

In more detail, it is worth recalling that the F-1 and F-2 flows have been defined as the superposition of a number of micro ON/OFF flows (see section 3 of the D6.4 report [1]). Each micro ON/OFF flow is directed from the metro network interconnection to a different access port.

So, when a micro flow is in OFF mode, no traffic is sent to a specific port. As specified in section 6.5 of the D6.4 report, after some seconds that a port is idle, the FPGA chip in the ETY1 and ETY2 devices shuts down the Rx-side of the SerDES interface connecting the destination access port. When the micro-flow transitions to ON mode, the status of the SerDES interface is returned in normal operation.

This operation permits to obtain a larger share of the energy saving reported in Fig. 74.

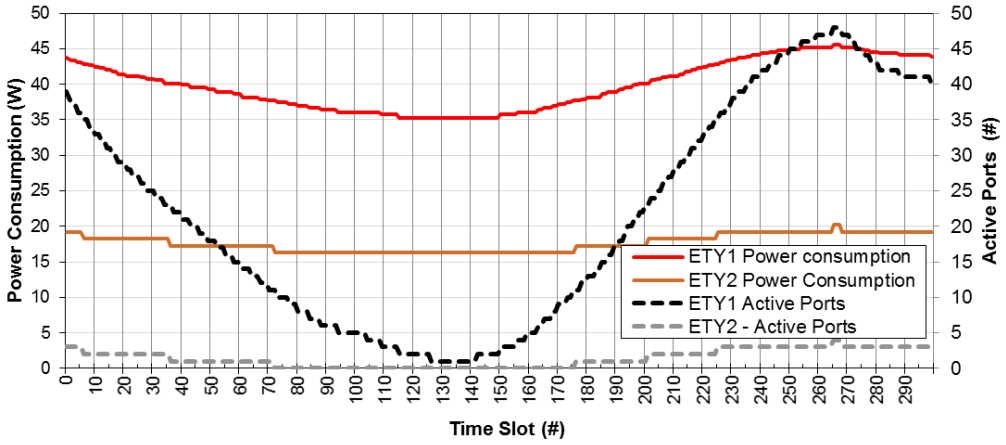

Fig. 74. Energy consumption and number of ports in active mode for the ETY1 and ETY2 prototypes in the access cloud.

3.2.6 Home Network Cloud

This section is organized as follows. In section 3.2.6.1, different aspects related to the home-gateway anatomy and to the origin of power consumption are discussed, with the support of some experimental results. Owing the achieved results, section 3.2.6.2 defines the configuration of the home gateways as they have been used in the main demonstration test. In section 3.2.6.3, the results of the main demonstration test are reported.

3.2.6.1 Testing setup and anatomy of energy consumption

In order to collect performance metrics, the power measurements were done with a single-shot power plug meter "Voltcraft SBC-500" with an accuracy of +/- 0.1W. The load was injected with the traffic generator (see D.6.4, Sect. 2 and 3).

As described in D3.3 the electrical current drawn by the home gateway is not constant, but fluctuates. Therefore, each single measurement result was obtained by averaging over a few (5…10) single-shot measurements. Table IX shows an example of this behaviour.

Table IX. Example of measurement samples on the power consumption of LQDE HG.

Measurement setup	Power [W]	Sample
Ethernet cable connected no Traffic, no WLAN RED home gateway	9.60	1
	9.30	2
	9.70	3
	9.60	4
	9.60	5
	9.60	6
	9.50	7
	9.40	8
	9.60	9
Average	**9.54**	
Standard deviation	**0.10**	

The LQDE EASY 80290 Box was evaluated in the three operating cases defined in D6.4: maximum power (RED), minimum power (GREEN) and in balanced operating mode (YELLOW).

It is worth recalling that the power consumption of the home gateway depends on:

- Configuration;
- Traffic load.

Configuration means the enablement of different functional blocks on the home gateway, while load refers to the data throughput. Power modulation by traffic load is small, while disabling unused functions saves a significant amount of energy.

The power management daemon of the YELLOW and GREEN home gateway not only adapts processor and DRAM frequencies, but also controls functional blocks:

- Automatic WLAN MIMO configuration depending on the load: 1x1 or 3x3,
- Automatic offloading of packet flows to the PPA packet processing accelerator,
- Automatic FXS power down if no analogue phone is connected,
- Automatic DECT ECO mode if no handset is associated,
- Automatic power-down of unconnected LAN, USB and analogue phone interfaces.

The PPA is a special processor, optimized for packet flow forwarding. If a high load packet stream persists over a certain time, it is offloaded by the MIPS host processor to the PPA. The load of the MIPS processor is then reduced and the clock frequency is adapted.

Disabled ports are scanned periodically to detect connected devices.

As shown in Table X, the power management daemon controls the frequencies of MIPS and DRAM in four steps.

Table X. Available frequencies of MIPS and DRAM components

Power State	CPU Clock	DDR Clock
D0	500 MHz	250 MHz
D1	393 MHz	196 MHz
D2	333 MHz	166 MHz
D3	125 MHz	125 MHz

The PMD daemon checks the CPU load against thresholds to increase or decrease the frequencies. For example, the VoIP subsystem requires a minimum CPU frequency of 200 MHz to operate. If the frequency is forced to 125 MHz, the voice quality degrades.

For RED and GREEN measurements, the power states have been fixed to D0 and D3, respectively. In addition, the automatic MIMO has been disabled and fixed to 3x3 and 1x1 MIMO, respectively.

For example, to assess the power consumption of an unconnected Ethernet port, these measurements were carried out:

- Router powered on, no device connected: (7.41 +/- 0.16) W,
- Ethernet cable connected to PC: (8.04 +/- 0.07) W.

The difference is the basic power consumption of an Ethernet interface: 0.63 W.

3.2.6.1.1 Power consumption depending on configuration

3.2.6.1.1.1 DSL

The DSL physical layer subsystem on the chip has constant power consumption. Over the DSL link there is always data transported. When in show time, framing, EOC and idle data are needed to maintain synchronization and data structure. Hence, the DSL subsystem in show time consumes a constant value of 0.89 W. This is also independent of the colour of the home gateway.

3.2.6.1.1.2 WLAN

The WLAN is very convenient, but contributes significantly to the home-gateway power consumption. Enabling both WLAN channels (2.4 GHz and 5 GHz) increases the power consumption of a free running home gateway by 4.41 W, from 7.24 W up to 11.65 W. Hence, 4.41 W is the base power consumption of the WLAN subsystem in the home gateway.

In the next step, data traffic of 50 Mb/s and 150 Mb/s has been transported over the WLAN interface. The results are in Fig. 75 for the RED home gateway. It does no automatic MIMO switching or other power saving measures, such as an on - chip bypass.

These measurements were made with local traffic sent from a LAN port to a tablet device via WLAN.

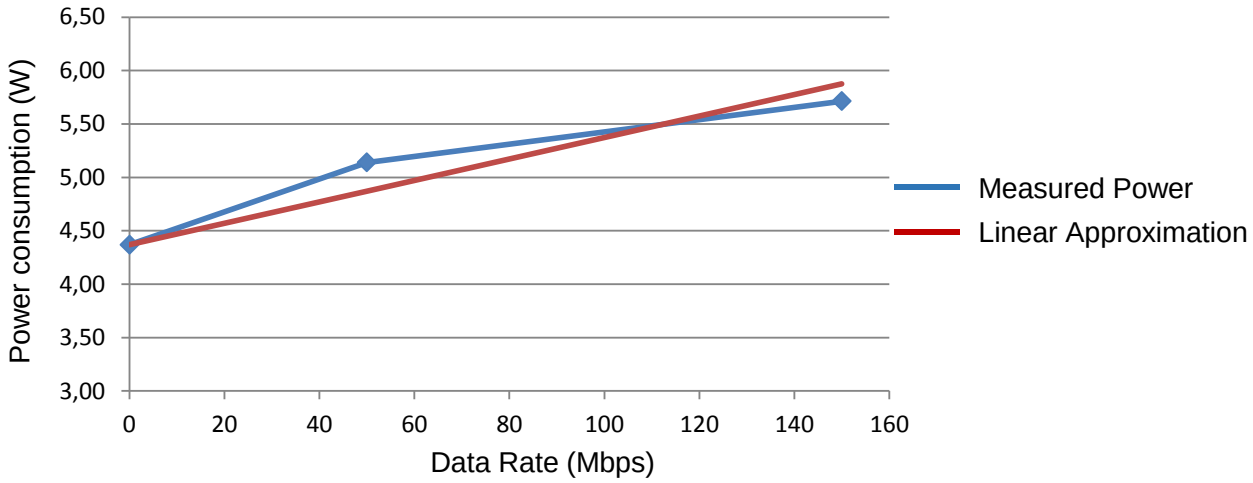

Fig. 75. WLAN power consumption and linear approximation according to the data rate.

Again, a best-fit linear approximation was made, leading to this formula:

$$P_{WLAN}(rate) = 4.41\ W + 0.01003\ W/Mbps\ (RED\ home\ gateway).$$

For the YELLOW home gateway, the zero bandwidth value is 3.14 W, due to automatic MIMO switch to 1x1.

3.2.6.1.1.3 Other Power Contributions

The DECT subset power contribution showed a constant value of 0.55 W, independent of whether a handset is associated or not.

Finally, the base power consumption of the home gateway was evaluated; it amounts to (7.41 +/- 0.16) W (see above). However, this value might be high compared to a commercial home gateway. The reason is that the "Blue Box" is not designed for energy efficiency, but for easy demonstration of all possible features of the chipset. For comparison, a commercial home gateway with the same interfaces was measured. Its free run (base) power consumption has been just a bit lower (7.19+/- 0.03) W.

3.2.6.1.2 Power Modulation

The load-dependent power modulation of the home gateway is small. This is depicted along with two measurements. For the first measurement, the two load values (a) no load and (b) 5 Mb/s were applied to a home gateway working in the following typical operation mode:

- VDSL line in operation with profile 17a
- One analogue phone connected, on-hook
- WLAN enabled, one client associated, low traffic
- Ethernet 100M interface with control traffic
- Ethernet 1000M interface with 0 or 5 Mb/s traffic from VDSL line
- Energy Efficient Ethernet (EEE) not working (since no remote EEE peer was available in the demonstrator setup – see the D6.4 report).

The results for the three colour home gateways are given in Table XI.

Table XI. Power consumption for the three available power states in idle mode (no traffic), and when a rate of 5 Mbps is received.

Mode	No traffic	Data rate 5 Mb/s
RED	11.01 +/- 0.04 W	11.34 +/- 0.22 W
YELLOW	9.98 +/- 0.15 W	11.05 +/- 0.20 W
GREEN	9.92 +/- 0.12 W	10.99 +/- 0.08 W

From the above table one can see that at zero load the RED home gateway wastes about 1 W compared to the YELLOW and the GREEN home gateways. The small difference between YELLOW and GREEN is expected. The uncertainty range of each measurement is given.

The same setup, using the RED home gateway, has been used for more extensive exploration. The RED home gateway has the advantage that no other power saving measures would influence the measurement. Traffic load was increased from zero to 120 Mbps, the maximum rate supported by the VDSL line. The resulting power consumption depending on the rate is shown in Fig. 76.

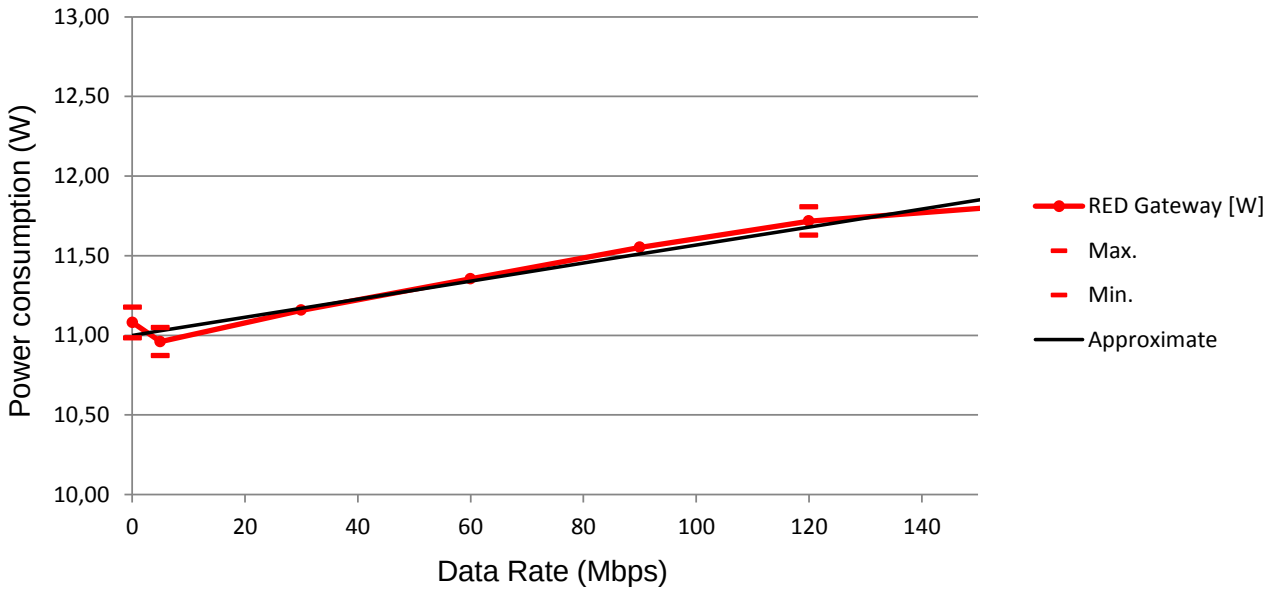

Fig. 76. Ethernet Test, RED home gateway averaged values.

In Fig. 76, an optimized linear approximation of the power consumption is shown. It is within the error ranges of the measurements – shown with three short horizontal bars – and simplifies further calculations. The formula is:

$$P_{LAN}(rate) = 0.63W + 0.00567\ W/Mbps.$$

Based on the linear approximation, the power modulation over a traffic profile similar to the ones used in the demonstration network has been calculated. For a better comparison, both load and power

consumption have been normalized to 100%, corresponding to 120 Mb/s and 11.68 W, respectively. The result is shown in Fig. 77. The message of this graph is that power dependency on the traffic load is limited. This statement is confirmed by technology. Interfaces such as Ethernet, WLAN, DECT consume a significant basic energy once they are enabled.

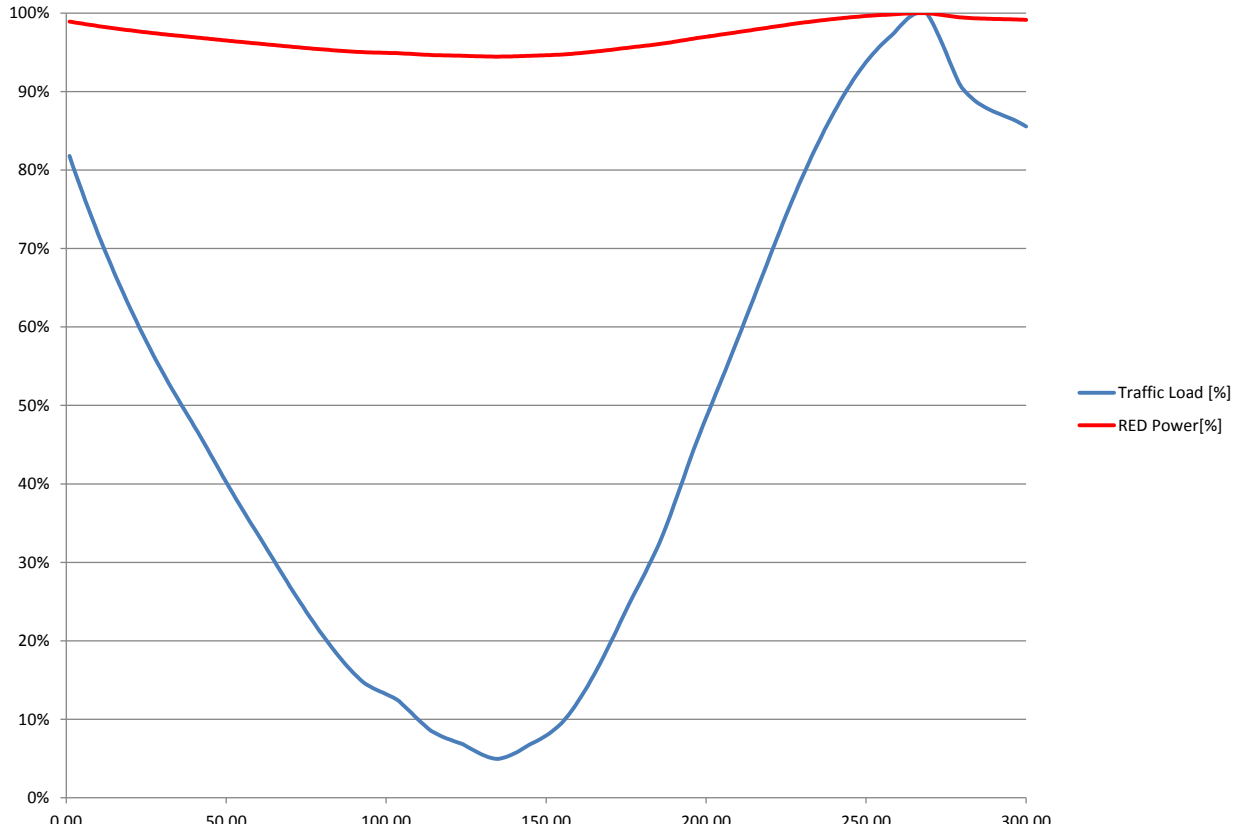

Fig. 77. Home gateway power consumption over ECONET profile F-2
x-Axis: time counts
y-Axis: load and power in %

3.2.6.2 Power and configuration modulation in typical use case

Power modulation by traffic load in the home gateway is small. To achieve significant savings, interfaces and/ or functional blocks must be switched off when not in use. These constant contributions have been evaluated in Table XII.

Table XII. Power consumption of main functional blocks.

Function	Base Power [W]
VDSL Line in Showtime	0.89 W
DECT enabled	0.55 W
Analogue phone connected (on hook)	0.08 W

Function	Base Power [W]
Analogue phone off hook	1.33 W

Switching off interfaces or functional blocks needs additional intelligence. Unused interfaces must be scanned periodically to detect new connected devices. Savings can be increased if the user is configuring down times of some of the functions. Comparison between WLAN and wired (non-EEE) Ethernet shows a clear advantage for wired.

For the typical use case a combined power modulation can be assumed, based on varying traffic load and dynamic switch-off and switch-on of modules.

Table XIII. Energy efficiency of Ethernet and WLAN.

	Base power	Gradient	Comment
Ethernet	0.63 W	0.00567 W/Mbps	Non-EEE
WLAN	3.14 W	0.01003 W/Mbps	

Table XIV. Maximum and minimum power configurations (energy consumption measures are referred to the RED state).

	Maximum Power	Minimum Power	Comment
Base power	7.41 W	7.41 W	
DSL Showtime	0.89 W	0.89 W	Always online
Ethernet	2 x 0.63 + 0.567 = 1.827 W	0.63 W	Daytime 2 ports with 100Mb/s traffic in the night 1 port without traffic
WLAN	3.14 + 0.50 = 3.64 W	0 W	Daytime 50 Mb/s Switched off in the night
DECT	0.55 W	0.55 W	Telephone always connected
Analogue phone	0.08 W	0.08 W	Always connected
Totals	***14.397 W***	***9.56 W***	***Maximum difference 4.837 W***

Therefore, during the demonstration, the three home gateways were set up with the "Minimum Power" configuration of Table XIV and the GREEN power state.

3.2.6.3 Demonstration Results

Fig. 78 and Fig. 79 show the experimental results collected during the main test in the ECONET demonstration in terms of aggregate energy consumption of the 3 green home gateways, and of traffic volumes arriving to them, respectively. Both results are expressed according to time slots.

As evident from the aforementioned figures, when the traffic volume drops to zero, the energy consumption of the home gateways slightly decreases, as expected (see Table XI).

Moreover, after three time slots (equivalent to approximately 15 minutes) from the traffic dropping, the three home-gateways were programmed to pass to the night configuration, as defined in the "Minimum Power" configuration of Table XIV.

Despite the small fluctuations in power consumption of Fig. 78, it can be noted that the savings with respect to the Business-as-Usual case span from a minimum value of 23% to a maximum value

of 40%. As explained in the previous subsection, this level of savings is mainly due to the combination of smart configuration settings and the internal power management of the HG.

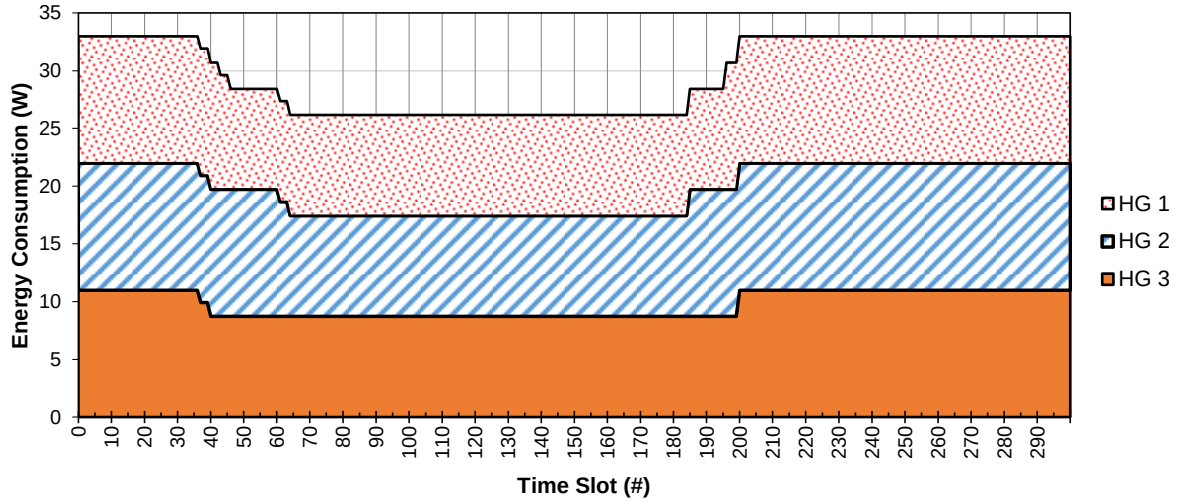

Fig. 78. Energy consumption of the 3 energy-aware HGs in the Home Cloud.

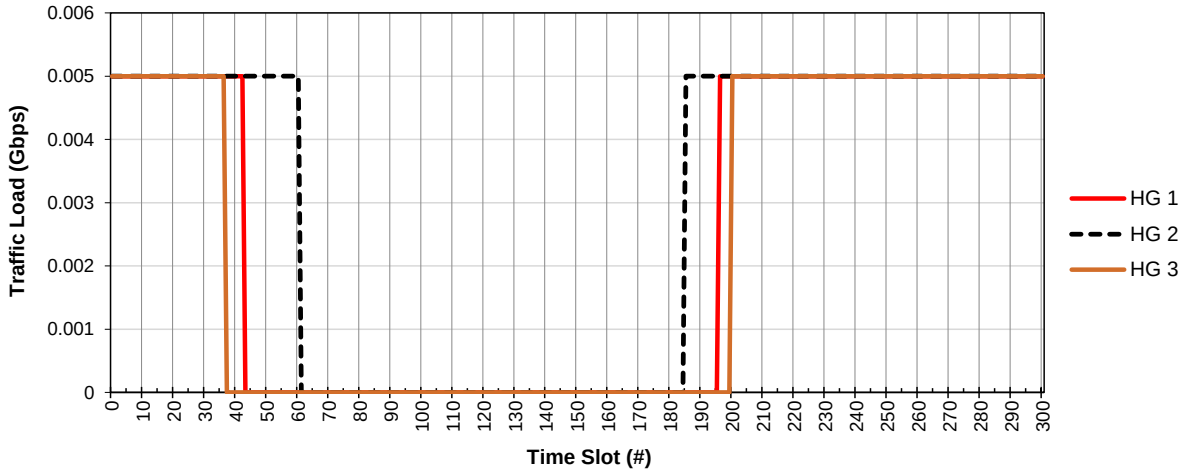

Fig. 79. Traffic volume of F-2 micro flows arriving at the 3 energy-aware HGs in the Home Cloud.

3.3 Performance evaluation of the Network Connection Proxy

The Network Connection Proxy (NCP) function, jointly developed by CNIT and INFOCOM, was tested in different scenarios to verify its correct operation and to collect performance measurements. Fig. 80 shows the integration of the NCP function into the ECONET demonstrator (see D 6.4 [1]). The NCP function is deployed in two customer sites, and a remote server is available in the core network. The presence of two customer sites accounts for two different deployments: in the first one the NCP runs on the home gateway, while in the second one the NCP runs on a device within the local network.

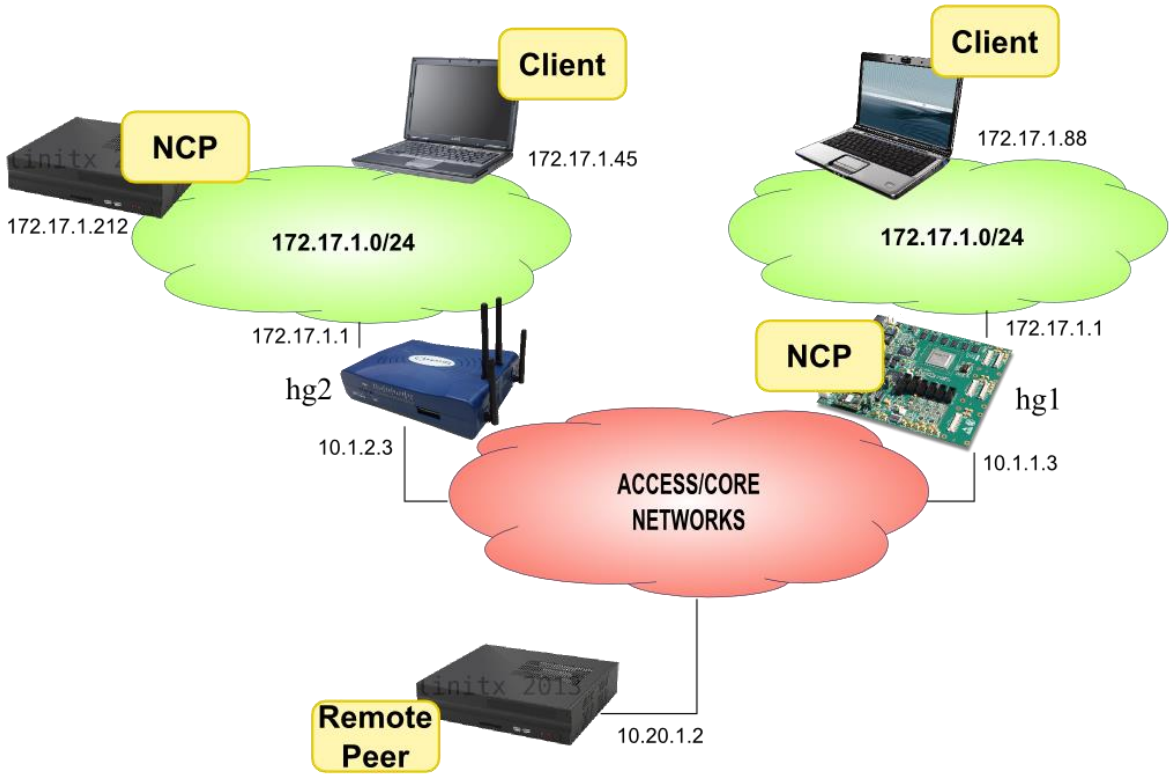

Fig. 80. Layout of the customer sites for benchmarking the NCP function in the ECONET demonstrator.

The evaluation of the NCP was aimed at the following:

- verifying the correct behaviour for all the actions defined in the project;
- carrying out a basic set of measurements to assess the impact of NCP operation on the user Quality of Experience (QoE).

3.3.1 Verification of the NCP behaviour

A set of basic operations has been defined in the project that covers most of basic networking operations an NCP could carry out on behalf of its clients. Table XV summarises the set of NCP operations defined so far, with explicit reference to their type and the relative action type (see the WP3 Internal Report [5] for more details about each type of operation). All the NCP operations can be registered by client devices through the UPnP interface, as defined in D 5.5.

Table XV. Summary of NCP operations defined so far.

		ARP	*Ping*	*DHCP*	*WoC*	*WoP*	*SRoP*	*TCP-KA*	*Hrt-Bt*
Type	*Network Connectivity*	✔	✔	✔				✔	
	Packet Management				✔	✔	✔	✔	
	Heart-beating			✔					✔
Action	*Wake up*				✔	✔		✔	✔
	Buffer packet							✔	✔
	Send predefined packet						✔		✔
	Build packet	✔	✔	✔				✔	
	Build packet by template								✔

The functional verification has thoroughly investigated the behaviour of each rule under the two different deployments depicted in Fig. 80. The analysis has included the presence of NAT operations performed by the Home Gateway device, and the issues that such network function implies. The main objective of the analysis was the following:

- the identification of bugs in the software implementation;
- the identification of issues that were not evident at the design stage.

As a matter of fact, the outcomes of the functional tests highlighted some criticalities in the original implementation that could not be considered in advance (see the WP3 Internal report [5] for full details) and that required some modifications of the NCP behaviour.

The most relevant of these outcomes are listed in the following.

- **Host unreachability errors**. When the NCP is supported in a networked device (as the Home Gateway, in our test bed), it was originally designed to "sniff" network packets in parallel to the device OS's packet handling chain. Hence, packets continue their trip once they crossed the NCP device and, inevitably, they trigger ICMP "host unreachable" errors by the device that should deliver them to their final destination. To fix this issue, it was necessary to make the NCP able to modify its hosting device's OS traffic handling, so to drop packets addressed towards sleeping hosts .
- **Spoofing of ARP caches in network switching equipment**. The NCP makes use of Gratuitous ARP to divert Ethernet packets when a client device changes its state (i.e. it either goes to sleep or wakes up); originally, the NCP sent Gratuitous ARP with the source MAC address of the sleeping devices, but it was evident that this puzzled the learning algorithm of Ethernet switches. All packets sent by the NCP on behalf of client devices must hence use the NCP's MAC as the source address.
- **Sequence of operations in transient periods**. The host's transition to and from the sleeping state requires some time; the NCP is notified of the transition slightly before the host goes to sleep and slightly after it resumes its operation, but it cannot exactly know when the device really stops or starts again using the network. That means there may be a

short transient where both NCP and the covered host are trying getting and answering the same packets. Finding the best way to perform all operations was quite complex, because the exact behaviour and all possible side effects must be taken into account.

- **Local address resolution**. Local ARP cache entries on the device hosting the NCP are not updated by responses sent by the NCP. The NCP must set static ARP entries by the system API.
- **Clients' Difficulty in notifying wake-up to the NCP**. A host attempt to establish a TCP connection with the NCP to notify its transition from the sleeping to the active state, happens when the NCP is still covering the host; thus, the ARP table of the NCP hosting device is altered, and no packets are sent back to the client host, hindering the conclusion of the three way handshake and the notification of the new state. To overcome this issue, early wake-up detection was developed that infers the current status of a covered host by its network activity.

3.3.2 Performance analysis

The prime target scenario for the NCP consists of a home network. The NCP is not expected to deal with a large amount of traffic; thanks to the basic filtering capabilities provided by suitable libraries, the number of processed packets should be quite limited, even when the NCP is traversed by all the traffic flowing between the site and the Internet (i.e., the NCP runs on the home gateway).

The most relevant performance indicators for the NCP are thus related to the goal of achieving "transparent" operation with respect to network hosts. The NCP hides the low-power state of the hosts it is covering, hence remote peer devices should not experience additional latency or service degradation with respect to the case when such hosts are active.

According to this, we focused our performance analysis on those NCP operations that most impact the user interaction. In this respect, the renewal of the DHCP lease, the answer to a TCP keep-alive message or the transmission of a heart-beating message impose no tight time constraint (a few seconds or less) for their execution, apart from the need to carry out the operation within the proper timeouts (generally, several hundreds of seconds, an internal that is well beyond the critical threshold for our NCP implementation). Instead, the answer to "ping" messages or the time to wake-up the host on the arrival of a new connection request directly affect the user; hence, it is worth focusing on these kinds of operations primarily.

3.3.2.1 Ping test

This test was conducted by analysing the delay to respond to ICMP echo-request packets, generated by the popular *ping* application. The relevance of this test was to assess the latency introduced by NCP operation on the round-trip time experienced by this traffic.

Table XVI shows the round trip time (RTT) and the packet loss measured under different conditions. The first two rows are measured when the target host is active, and are reported for comparison only. The following rows were collected while the target host was sleeping, thus the NCP answered the echo-request messages on its behalf. We ran the test in different conditions, by varying the frequency of the echo-request messages and by changing the number of devices served by the NCP. As one can see, the RTT is quite a lot larger when the NCP is involved; this is due to sniffing operations, which delay the delivery of packets to the NCP for processing; however, the absolute

value of the RTT is small and does not significantly affect the user. Further, note that the relative errors will become negligible once *ping* commands from the Internet were considered.

It is also important to notice that no packet loss occurs. To further investigate the possibility of packet loss during transient periods, we repeated the measurement while making the NCP-covered host switch from the active mode to standby, while the ping session was active: the result of the test is plotted in Fig. 81. We can see the average values for the RTT are congruent with Table XVI; furthermore, we verified that no packet loss occurs.

Table XVI. Performance measured on "ping" traffic without and with the intervention of the NCP. Different trials were collected by changing the time interval of the echo-request packets and the number of devices served by the NCP. Measures include statistics about the Round Trip Time and the packet loss ratio. Results for each trial were averaged over 100 samples.

	# of NCP clients	ICMP packet interval	Round Trip Time (ms)				Packet loss (%)
			Min	Average	Max	Mean dev	
Active host (w/o NCP)	-	1 sec	0.243	0.380	0.514	0.054	0
Active host (w/o NCP)	-	500 ms	0.142	0.215	0.341	0.046	0
Sleeping host (w NCP)	1	1 sec	1.134	1.407	1.734	0.140	0
Sleeping host (w NCP)	1	500 ms	1.171	1.388	1.748	0.141	0
Sleeping host (w NCP)	1	200 ms	1.181	1.408	1.816	0.139	0
Sleeping host (w NCP)	10	500 ms	1.151	1.399	1.699	0.150	0
Sleeping host (w NCP)	10	200 ms	1.157	1.391	1.760	0.137	0

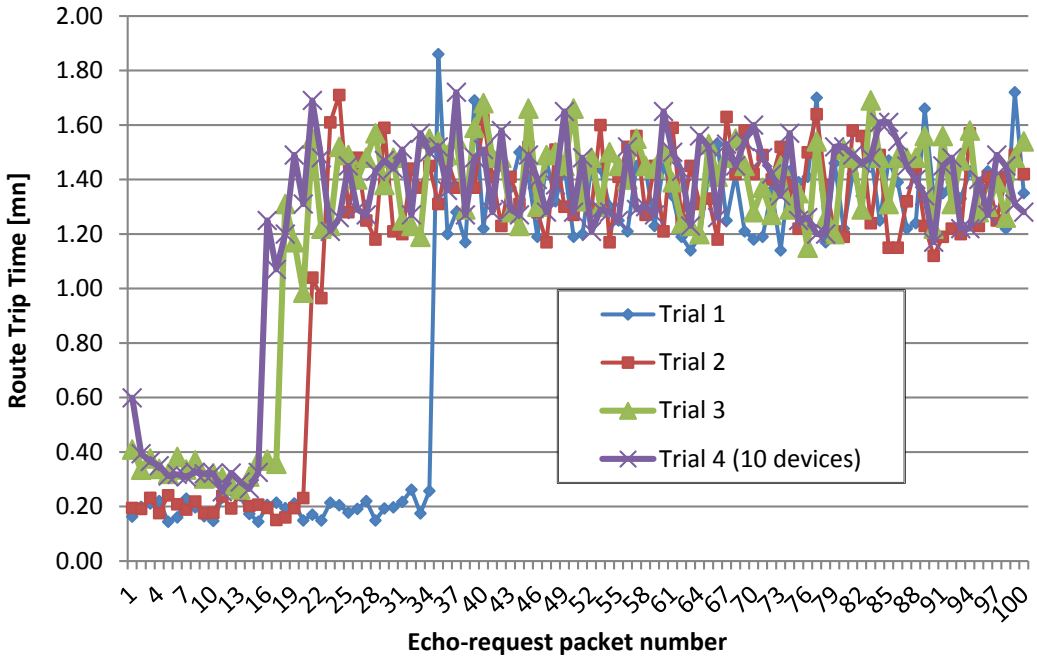

Fig. 81. Round Trip Time measured during the switch of the target device from active mode to standby. Two different trials are reported, the switching times are not synchronized.

3.3.2.2 Wake on connection test

This test assesses the latency experienced by the user when trying to access a service on the sleeping host. Here, we assume the service must be directly provided by the target host, so wake-up is required; this behaviour is implemented by the NCP through the Wake-on-Connection rule (see the Internal Report [5]).

We used SSH as the target application. Our analysis considers the responsiveness of the NCP to process the SSH request and the time the first packet from the target host gets back to the SSH client. Results from our trials are shown in Table XVII.

Table XVII. Summary of performance evaluation for new connection requests. An SSH client was used to initiate new requests towards a sleeping SSH server (protocol TCP, port 22). The analysis considers the time taken by the NCP to send out the first WOL packet to the sleeping device and the latency to receive the answer to the TCP SYN packet from the SSH client. Several trials were performed, changing the number of devices registered at the NCP.

	# of NCP clients	Statistics over 10 trials				Average # of attempts
		Min	Average	Max	Std. dev.	
WOL request	1	1.549 ms	1.747 ms	1.894 ms	0.123 ms	3.3
	10	1.594 ms	1.885 ms	2.148 ms	0.176 ms	3.1
Answer to the TCP SYN request	1	3.67 s	5.71 s	7.99 s	1.623 s	3.3
	10	3.23 s	5.13 s	7.62 s	1.34 s	3.1

Our evaluation showed the latency due to the NCP is almost negligible (just a few milliseconds on average), while most of the latency comes from getting an answer from the sleeping device (about 5 seconds on average). The latter is ascribable to two factors:

a) the time to wake up the host, which approximately contributes for 3-4 seconds;

b) the time to detect the device is awake, which is rather variable and accounts for 3-4 seconds as well.

We found that once the host is awake, it does not start immediately to send out network traffic, although it receives packets from the network; the execution of the same application in charge of notifying the host power state to the NCP device is delayed by several seconds. The latency in setting up an SSH connection towards a sleeping device is noticeable by the user, yet the figures shown in Table XVII can be considered acceptable for most users[1].

3.3.2.3 Heart-beating and TCP migration

The NCP sends periodic unsolicited or solicited heart-beating on behalf of client devices, both over UDP and TCP connections. In this perspective, assuming the typical time intervals are within the range of hundreds of seconds or tenths of minutes, performance is not a matter in this case. However, heart-beating over TCP requires transparent and seamless migration of the session between the client device and the NCP; this operation must complete before the client falls into standby. Hence, we measured the time taken to suspend the TCP session on the client device and to transfer its current state to the NCP (see Table XVIII).

Table XVIII. Performance analysis for TCP migration. Latency measured includes the time to freeze the TCP session, to save its status and to transfer the status to the NCP through the UPnP interface.

	Min	Average	Max	Std. dev.
1 NCP client registered	128 ms	131 ms	133 ms	1.58 ms
10 NCP clients registered	128 ms	132 ms	136 ms	0.24 ms

Figures shown in Table XVIII are small, yet they must be related to the mechanism implement to go to standby. Indeed, if applications are simply notified once this process is initiated, the above values are not sufficient to guarantee the procedure completes before the network interface and the host are shut down; instead, if applications could slow down such process, there would be no problems. Such issue is still under investigation in the NCP implementation.

3.3.2.4 Registration of actions

The registration of packet processing actions needed by each client while it is in the standby state happens through the UPnP interface. We measured the time this operation requires; results are shown in Table XIX.

[1] Note the figures shown in Table XVII are indeed comparable with values over slow Internet links.

Table XIX. Delays introduced by the registration procedure. All figures are in milliseconds. Values are relative to the time taken by an application to execute all the process of registering an action (i.e., call to local internal API to register actions by UPnP). Results were averaged over 10 trials for each action.

	Registration (UPnP+NCP processing)				NCP processing			
	Min	Average	Max	Std dev	Min	Average	Max	Std dev
Ping	105.914	106.190	106.488	0.180	0.912	0.955	1.031	0.043
WakeOnConnection	105.250	105.336	105.396	0.046	0.812	0.844	0.912	0.031
ProxyDHCP	105.168	105.337	105.584	0.149	0.832	0.859	0.906	0.026
WakeOnPacket	105.254	105.389	105.544	0.082	0.820	0.837	0.884	0.023
TCPKeepAlive	102.665	104.924	105.642	1.193	0.580	0.797	0.888	0.117
Heartbeating	102.740	105.189	105.862	0.921	0.833	0.905	1.274	0.131
All traffic	102.665	105.394	106.488	0.735	0.580	0.866	1.274	0.092

Results show that the registration procedure takes about one hundred milliseconds, independently of the action type; further, such values are very stable across different actions and trials. However, note that most of the time is taken by the UPnP stuff, as the delay ascription of NCP processing of data is negligible.

3.3.3 Overhead

UPnP messages are exchanged to advertise presence in the network, broadcast discovery messages at periodic intervals, download the device XML description file during registration with the Control Point (CP), periodic device subscription renewal messages, etc. In this Section, network overhead due to UPnP signalling is evaluated under steady state conditions, action registration and power state notification.

We performed an experiment that lasted almost 12 minutes. The overhead in terms of number of packets and bytes exchanged is shown in Fig. 82; a brief description of the events is given below:

1. Start of client application. During this event, both the control point and the controlled device on the client advertised their presence by broadcasting a certain number of SSDP messages.
2. Start of NCP application and client registration with NCP. NCP CP and controlled device broadcasted their presence. The number of packets and bytes transferred are quite large, as the control point at both ends downloaded the device description file (XML) from the URL provided during the SSDP discovery phase. Furthermore, for multiple received SSDP discovery messages, CPs on both ends downloaded device description files in order to know if the device is already registered or not.
3. Steady state condition. This phase only contains the UPnP heartbeats: the NCP continuously renewed its subscription to the state variables it is interested in[2].
4. Action Registration. The client registered a WakeOnConnection action at the NCP.
5. Client power state change notification to the NCP.

[2]The subscription renewal timeout was set to 10 seconds in the experiment; the subscription is renewed at half of this value.

6. During this period, the client was in sleep mode.
7. Host wakeup. The client informed the NCP about the power state change and the NCP stopped covering for that device. Both the NCP and the client exchanged notifications of changes in the state variables.
8. Host de-registration from NCP. This was the bye-bye message sent by the NCP to the client CP. The client CP received multiple bye-bye messages and each time it downloaded the XML description file of the NCP device to check if the device was in the registered devices list or not.
9. Client bye-bye message broadcasted at client application exit.

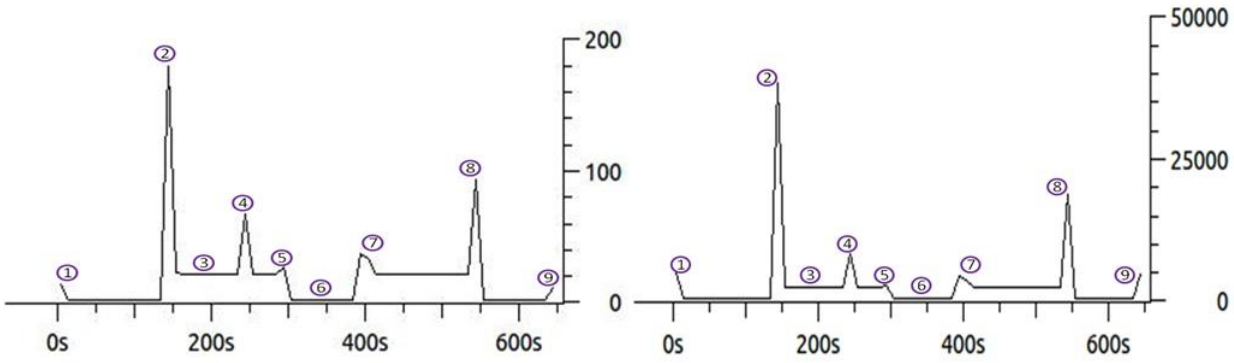

a) Number of packets vs Time b) Bytes exchanged vs Time

Fig. 82. UPnP network overhead.

The breakdown of the number of packets and bytes exchanged for different events are shown in Table XX. From Table XX, it is clear that most of the packets were exchanged during the steady state condition, while the average packet size was the smallest. The average packet size during the discovery and description phases is the largest, as the CPs at both ends downloaded the device description XML file to discover the device details.

Table XX. Breakdown of packets exchanged during the experiment.

Event Number	Event Type	Number of Packets Exchanged	Total Bytes Exchanged	Avg. size of a packet [B]
1,2	Discovery & Description	130	34678	266.75
3	Steady State	650	68154	104.85
4	Action Registration	46	5976	129.91
5,7	Power State Notification	43	4929	114.62
8,9	UPnP un-subscription	80	17508	218.85

Fig. 83 and Fig. 84 show the percentage of bytes and packets exchanged during different events, respectively. Fig. 83 shows that 52% of bytes were exchanged during the steady state period, whereas the discovery and description phases contribute to 26%. Similarly, Fig. 84 shows that most of the packets were exchanged during steady state condition (68%) and the discovery phase contributed to

14%. It can also be observed that the overhead during action registration and power state notification was quite low compared to the steady state and the discovery phases.

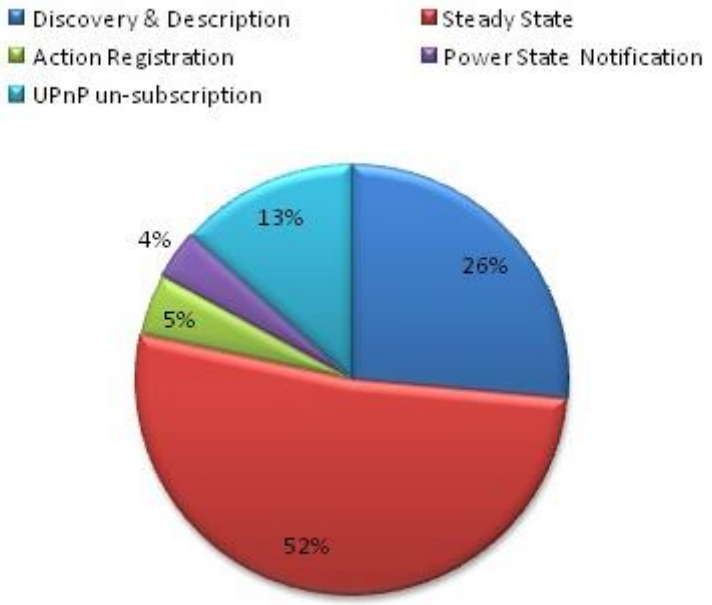

Fig. 83. Percentage bytes exchanged during different events.

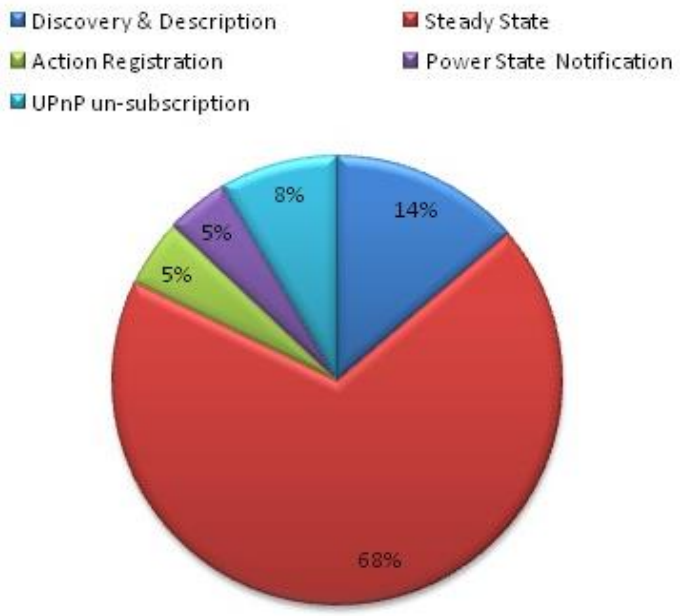

Fig. 84. Percentage number of packets exchanged during different events.

4 The expected impact of ECONET technologies

As already sketched in Section 3.1, the results of the demonstrator are not fully representative of the impact of the ECONET technologies on real network infrastructures for the following main reasons:

- It was not possible to provide *typical densities of devices* usually deployed in network segments of Telecom Operator infrastructures.
- The ECONET prototypes have been completed with *"legacy" components* (without green mechanisms).
- *The dimensioning of the links and of the devices with respect to the traffic load* is not in line with the one adopted in real network deployment.
- Some of the *prototypes do not include all the possible energy-saving solutions* studied in the project.

For these reasons, the goal of this section is to provide an accurate estimate of the potential impact of the ECONET technologies when massively applied to a real network infrastructure of telecommunications operators.

In order to perform this analysis, it is necessary to define an analytical framework able to represent the experimental behaviour of ECONET prototypes during the demonstration, and to synthetically provide their energy profiles according to the incoming traffic volumes. For simplifying this analysis, but without losing generality, energy-aware primitives were not modelled in detail, but simply grouped into two representative sets; namely, standby and Dynamic Adaptation (DA). The fine-grained features and characteristics of each green primitive used in the demonstration were taken into account by means of a careful parameterization of the model.

It is worth noting that the analytical framework here proposed, as well as the reference scenarios and the data used, directly come from the extensions of the one used during the project proposal [7], and of the refined one developed in the context of WP2 for estimating the potential impact of the project [8][9]. This choice permits a direct comparison between the targets declared at the project start, and the impact estimated from the experimental results of the final demonstration.

Like the model in [8], the energy consumption of the "inertial" improvements in network technologies are not considered. Only the effect of green technologies and solutions at data- and control-planes are explicitly modelled in the analytical framework.

The section is organised as follows. Subsection 4.1 details the energy profile model for representing the energy-aware states available in network devices, while Subsection 4.2 briefly presents the network infrastructure and the three reference scenarios. Finally, Subsections 4.3 and 4.4 report the results of this estimate by putting emphasis on the energy consumption and on the relative energy savings in the different ECONET demonstrator clouds.

4.1 The Energy Profile Model

This section introduces the analytical framework for modelling the impact of energy adaptive technologies and solutions adopted in the ECONET demonstrator prototypes. It aims to extrapolate the performance of these prototypes when used in operating scenarios that are different from the demonstration one.

The entire model is designed to capture the main effects and benefits of the energy saving primitives in a highly parametric way, without entering into details on how these primitives are realized inside device platforms. We used an extended set of efficiency parameters, which allows us to represent all the prototypes.

For this purpose, the model is based on the "*energy profile curve*" $\phi_{dyn}(\lambda)$ of network devices. This curve is defined in a similar way with respect to the GAL autonomic states: it represents how the power consumption of a device changes according to the incoming offered load λ. In more detail, a curve for each ECONET demonstrator cloud has been defined to represent the average behaviour of a device working on that network segment. The $\phi_{dyn}(\lambda)$ curves have been obtained by fitting the energy consumption data collected during the demonstration (see Fig. 12 at page 18 of this report). Since the data range does not cover all the offered load values, it was decided to maintain the minimum power consumption constant over all λ values between 0 and the minimum load value measured in the demonstration.

Further, two parameters are needed by the analytical framework to completely characterize the ECONET impact, namely:

- ϕ_{static} – static energy consumption of the device (or link) consumed even if it is idle (i.e., no packets received or transmitted).
- ϕ_{sleep} - energy consumption when the device (or link) is in a low consumption standby state.

A device is supposed to potentially provide N_{P-PsS} Service rates, one for each available Primitive Power Scaling Substate (P-PsS). For the sake of simplicity, these rates have been modelled as a discrete set uniformly distributed in the range $[0 - \mu_{max}]$, where μ_{max} represents the BAU device capacity. The n^{th} element of this set is computed in the following way:

$$\mu_n = \frac{\mu_{max}}{N_{P-PsS}} \cdot n \tag{1}$$

where $n = 1, \dots, N_{P-PsS}$.

The service rate set is ordered in n: $n = 1$ corresponds to the state with the lowest energy consumption and forwarding capacity; $n = N_{P-PsS}$ is the most energy-hungry state that provides the highest network performance. In more detail, the N_{P-PsS} state corresponds to the BAU case.

The energy-aware device selects the minimum service rate that supports the incoming traffic λ:

$$\mu(\lambda) = \min_{i=1,\dots,N_{P-PsS}} \{\mu_i | \mu_i \geq \lambda\} \tag{2}$$

When a device does not perform any operations, it can be put in idle or standby state and it consumes ϕ_{sleep}. The time spent in such state mainly depends on the average traffic utilization, as $T_{idle} = 1 - \frac{\lambda}{\mu(\lambda)}$

Thus, the average consumption of a device working in a generic cloud can be expressed as follows:

$$\phi(\lambda) = \phi_{static} + \left(\phi_{sleep}^{\frac{1}{\sigma}} \cdot \left(1 - \frac{\lambda}{\mu(\lambda)}\right) + \phi_{dyn}^{\frac{1}{\sigma}}(\lambda) \cdot \frac{\lambda}{\mu(\lambda)} \right)^{\sigma} \tag{3}$$

where σ is the shape parameter of the power consumption curve. In detail, the case $\sigma = 1$ corresponds to the ideal case, where there is no overhead in entering and exiting from standby/idle states and the energy consumption increases linearly with respect to the traffic load λ. We assume $\sigma < 1$ for real implementations of idle/standby logic [11].

As previously sketched, the analysis of the project impact needs also to consider:

- The maturity level of prototypes, or the fact that they only include a part of the developable green mechanism proposed by ECONET;
- The share of legacy hardware used in the demonstrator.

This aspect has been handled by introducing the parameters ψ_{static}, ψ_{dyn}, ψ_{sleep} and $\xi_{static-dyn}$ that represent the improvement due to a better maturity levels of prototypes, and the absence of legacy hardware:

- ψ_{static} – percentage of energy reduction of ϕ_{static} considering the technological maturity of the energy-aware primitives.
- ψ_{dyn} – percentage of energy reduction of ϕ_{dyn} considering the technological maturity of the energy-aware primitives.
- ψ_{sleep} – percentage of energy reduction of ϕ_{sleep} considering the technological maturity of the energy-aware primitives.
- $\xi_{static-dyn}$ – percentage of ϕ_{static} energy consumption that passes to ϕ_{dyn} due to the absence of legacy hardware.

Considering Eq. 3 and the ψ_* and $\xi_{static-dyn}$ parameters, the average energy consumption of a "well-engineered" green device in an "all-green" network can be expressed as:

$$\tilde{\phi}(\lambda) = \tilde{\phi}_{static} + \left(\tilde{\phi}_{sleep}^{\frac{1}{\sigma}} \cdot \left(1 - \frac{\lambda}{\mu(\lambda)}\right) + \tilde{\phi}_{dyn}^{\frac{1}{\sigma}}(\lambda) \cdot \frac{\lambda}{\mu(\lambda)} \right)^{\sigma} \quad (4)$$

where:

$$\tilde{\phi}_{static} = \phi_{static} \cdot (1 - \psi_{static}) \cdot \left(1 - \xi_{static-dyn}\right) \quad (5)$$

$$\tilde{\phi}_{dyn}(\lambda) = \frac{\phi_{BAU} - \tilde{\phi}_{static}}{\phi_{BAU} - \phi_{static}} \cdot \phi_{dyn}(\lambda) \cdot (1 - \psi_{dyn}) \quad (6)$$

$$\tilde{\phi}_{sleep} = \phi_{sleep} \cdot (1 - \psi_{sleep}) \quad (7)$$

Regarding the traffic load to be used in the estimation, it was decided to adopt the same profiles used in the demonstration – including the "day & night" shaping curve (see section 3 of the D6.4 report) – but scaling the peak load of flows in order to assure the common level of *overprovisioning* of network resources.

In fact, while the peak utilization of links and devices in the demonstrator stayed, in some cases, around 60%-90% of their capacities, in real network infrastructures the utilized bandwidth never exceeds 50% of the deployed capacities. In more detail, network capacity is usually dimensioned, or

better "overprovisioned", on the peak traffic load estimates over a lifetime period of deployed devices in order to never exceed 50% of the capacity.

For this reason, we define the parameter ξ that represents the degree of network capacity *overprovisioning*. In more detail, the role of this parameter is to scale the traffic level that was used during the demonstration experiments in order to achieve the desired degree of network overprovisioning.

4.2 Network Infrastructures and Scenarios

In order to perform a fair comparison with the potential impact estimate of the ECONET technologies, which was included in the project "Description of the Work" and published in prestigious international journals and magazines [7] [8], it has been decided to rely this final analysis on a reference network infrastructure as close as possible to the one used in such articles.

In more detail, as defined in the D6.4 report, the network is supposed to be organized in six "clouds", which represent different network segments namely: *home*, *access*, *datacenter*, *metro*, *core* and *transport*.

For each cloud, the density and the average energy consumption of devices have been fixed, where possible, to the same levels defined in references [7] and [8]. In this respect, it is worth noting that two further clouds are considered.

One of these additional clouds comes from the splitting of the metro and transport segments, which in the initial estimate were considered part of the same cloud. However, the number of devices and related average energy consumption are kept homogenous with respect to the starting network scenario.

The second additional cloud is represented by the datacenter network, which was not targeted explicitly at the beginning of the project. In this case, 2000 switches with an average energy consumption of 1.2 kW have been included in the reference network. It is important to underline that, in order to perform a meaningful and fair comparison with the initial impact estimates, the device density in the datacenter cloud has been heavily underestimated. So, the real potential impact of the ECONET technologies in this network area may be much broader than the levels here estimated.

As shown in Table XXI, we defined for the k^{th} cloud the device density through the $N_d(k)$ parameter, describing the potential network infrastructure that TELIT is going to set up for the period 2015-2020.

Table XXI. Device density and average energy consumption for the home, access, metro, transport, core and datacenter network segments of the reference infrastructures defined in [7], [8], and [9].

Network segments	k	$\omega_c^{(k)}$	$\widehat{\phi}_c^{(k)}$ [W]
Home	1	17,500,000	10
Access	2	27,344	1,280
Metro	3	1,250	6,000
Transport	4	500	6,000
Core	5	175	10,000
Datacenter	6	2,000	1,200

4.3 Cloud Power Consumption and Energy Saving

This section introduces the main key performance indexes used to estimate the ECONET impact in the three scenarios introduced in Section 4.2.

The total power consumption of the k^{th} cloud can be expressed as in the following:

$$\Phi_c^{(k)} = \frac{\omega_c^{(k)}}{N_{slot}} \sum_{i=1}^{N_{slot}} \tilde{\phi}^{(k)}(\lambda_i) \tag{8}$$

where N_{slot} is the number of time slots in a day, $\omega_c^{(k)}$ is the number of network devices of the k^{th} cloud, as defined in Table XXI, and $\tilde{\phi}^{(k)}(\lambda)$ is the average energy consumption of a device in the k^{th} cloud when the average offered traffic load is λ. The power consumption in the BAU scenario can obviously be obtained by replacing $\tilde{\phi}^{(k)}(\lambda)$ with ϕ_{BAU}, and it turns out to be $\Phi_{BAU\ c}^{(k)} = \omega_c^{(k)} \phi_{BAU}$.

The energy saving over the horizon of a year achievable in the k^{th} cloud by the adoption of ECONET technologies against the BAU scenario can be simply expressed as follows:

$$\Delta E_c^{(k)} = \left(\Phi_{BAU\ c}^{(k)} - \Phi_c^{(k)}\right) \cdot \Delta y \tag{9}$$

where Δy represents the number of hours in a year, and $\Delta E_c^{(k)}$ is meant to be expressed in GWh per year.

Our impact evaluation of the ECONET technologies is completed with two further indicators, expressing the impact in terms of the equivalent of tCO_2e (tons of Equivalent carbon dioxide) and Euro (€) saved during one year of exercise of the ECONET scenario against the BAU one.

The equivalent of tCO_2e saved during one year can be computed as follows:

$$\eta_{tCO_2e}(k) = \Delta E_c^{(k)} \cdot \epsilon_{tCO_2e} \tag{10}$$

where $\eta_{tCO_2e}(k)$ is the equivalent of tCO_2e saved during one year with respect to the BAU scenario, and $\epsilon_{tCO_2e} = 0.7$ tCO_2e/Wh[3] is the quantity of tCO_2e consumed with 1 Wh of energy.

Similarly, the amount of € saving in one year can be computed with Eq. 11:

$$\eta_{€}(k) = \Delta E_c^{(k)} \cdot \epsilon_{USD} \tag{11}$$

where $\epsilon_{€} = 0.2$ €/Wh[4] is the cost in € for 1 Wh of electric energy.

[3] Source: http://www.carbonneutralcalculator.com/Carbon%20Offset%20Factors.pdf
[4] Source: http://epp.eurostat.ec.europa.eu/statistics_explained/images/c/cf/Electricity_prices_%28EUR_kwh%29_%28NAT_kWh%29_2013s1.png

4.4 Energy Saving Gain Results

In this subsection, the estimation results are shown. The reference network introduced in section 4.2 is considered. The traffic levels correspond to the ones shown in Figs 16-39, scaled according to the overprovisioning degree ξ. Just like in the demonstration setup, the traffic flows used in the estimation follow the same "night & day" profile divided into 300 time slots.

Three different scenarios have been considered in the present estimation, namely:

Green Demonstration scenario, where the network devices provide energy-aware functionalities as the ones deployed during the ECONET project (and therefore without applying the improvements in Eqs. 4-7). In such case, the raw experimental results have been elaborated only to scale the traffic to the defined overprovisioning levels, and suitably weighted for representing the typical device densities in the different network segments (see Table XXI). The share of legacy hardware is kept at the same value with respect to the demonstrator. The applied overprovisioning levels are the ones reported in the bottom line of Table XXII.

- *Short-term Impact* scenario, obtained in the same conditions of the scenario above, but considering a network infrastructure composed only by fully green devices (i.e., no legacy cards or modules inside green devices) with the same mechanisms integrated into the demonstrator. To reflect these aspects, Eqs. 4-7 were applied by keeping all the parameter values fixed to 0, except for $\xi_{static-dyn}$ that has been dimensioned for removing the energy consumption bias due to legacy hardware (see Table XXII).
- *Long-term Impact* scenario, as in the previous case, the network is entirely composed of fully green devices, but a technological improvement/refinement is considered by applying Eqs. 4-7 with the parameter values in Table XXIII. These technological improvements are meant to come from a better maturity of device platforms and of green mechanisms, as well as from the integration of power saving solutions studied and analysed by ECONET, but not included into the demonstrator.

As previously sketched, the *Business-As-Usual* (BAU) scenario, where the network devices do not any provide energy-aware functionalities, is taken as a term of comparison.

Additional details on the preliminary assumptions used to produce this impact estimate and on the use of raw experimental results have been reported in Annex I of this deliverable report. This Annex also includes a comprehensive outlook on all the operations made to derive the impact estimate in all the considered scenarios. For all of them, we obtained the values of the key performance indexes defined in section 4.3, namely, the average energy savings, the cut of operating expenses, and the reduction of greenhouse gas emissions. All these values have been reported and expressed both for each network segment and for the entire network infrastructures.

Table XXII. Parameter setting for the short-term impact estimate in all the cloud networks.

Parameter	Transport	Core	Metro	Datacenter	Access	Home
σ	0.90	0.75	0.75	0.75	0.90	0.90
N_{EAS}	2	10	12	2	3	3
ψ_{static}	0%	0%	0%	0%	0%	0%
ψ_{dyn}	0%	0%	0%	0%	0%	0%
ψ_{sleep}	0%	0%	0%	0%	0%	0%
$\xi_{static-dyn}$	40%	35%	40%	30%	0%	25%
ξ	50%	50%	50%	50%	50%	30%

Table XXIII. Parameter setting for the long-term impact estimate in all the cloud networks.

Parameter	*Transport*	*Core*	*Metro*	*Datacenter*	*Access*	*Home*
σ	0.90	0.75	0.75	0.75	0.75	0.80
N_{EAS}	5	10	12	4	5	5
ψ_{static}	30%	20%	47%	30%	55%	45%
ψ_{dyn}	25%	5%	33%	10%	23%	25%
ψ_{sleep}	30%	5%	45%	7%	18%	7%
$\xi_{static-dyn}$	50%	45%	50%	34%	15%	50%
ξ	50%	50%	50%	50%	50%	30%

Figs. 85, 86, 87 report the estimated values of the three key performance indexes for the Green Demonstration scenario. In this respect, we underline that the energy saving levels for each cloud correspond to the ones obtained with the demonstrator and the assumptions reported in the Annex. These saving percentages have been applied to the device densities and average power requirements of the reference network infrastructure, introduced in Section 4.2 and in Table XXI, in order to obtain the energy saving of the entire network, and the related impact in terms both of Green House Gas emissions and of OPEX. These operations have been repeated also for the scenarios analysed in the following.

The obtained results outline an impressive overall energy savings of 39.7%, corresponding to a cut of operating expenses (OPEX) of approximately 164 M€ in a year, and a reduction of greenhouse gas emissions of 575 $MtCO_2e$ (equivalent to the carbon footprint produced by 58,000 cars in a year). It is worth noting that:

- These indexes give an overall impact of the ECONET technologies quite close to the project short-term target (50% of energy savings), notwithstanding the large share of legacy hardware used in the demonstrator. This is a clear proof of how the proposed technologies can have a significant impact also when applied in few devices/cards in the network.
- A share of more than 20% of the aforementioned indexes is due directly to the operator infrastructure, while the remainder derives from the end-user homes.

All the network segments exhibit significant average savings:

- from 26% in the transport segment, where power management techniques are not applied on the current traffic volumes, but on pre-determined reservation levels, and consequently they turn to be more conservative than the ones used in other clouds;
- to 71.6% in the core cloud, where parallel hardware performing complex IP-level operations can be suitably managed during the day to save a huge amount of energy.

The other segments exhibit average savings spanning in the range of 33%-40%, and the same comments to raw demonstration results (see section 3.2) can be extended here.

The Short-term Impact results are shown in Figs. 88, 89, and 90. In this case, an average energy saving of 51.5% is obtained across the entire network chain, which is a saving that even exceeds the short-term target of the project, and 11.8% more than in the previous scenario. In terms of OPEX cut, the massive adoption of ECONET technologies may result in a saving of more than 200 M€ per year, avoiding emissions for approximately 700 $MtCO_2e$, which correspond to the carbon footprint of 70,000 cars in a year.

It can be underlined that, as expected, the increase of savings is not linearly proportional to the share of ECONET devices deployed in the network. In fact, by removing the legacy hardware bias for the $\xi_{static-dyn}$ shares shown in Table XXII (more than 30% in most of the clouds), the overall energy saving of the entire network infrastructure increases only about 12%. This effect is clearly due to the fact that, even if the share of ECONET devices enlarges, the green capabilities of some of these devices cannot be fully exploited, since they need in any case to remain active at capacities close to the BAU one in order to serve the incoming traffic load with a suitable QoS level. More detailed comments can be found in the Annex document.

Figs. 91, 92, and 93 show the results for the Long-term Impact estimate of the ECONET technologies.

These results are obtained with the parameters showed in Table XXIII, which were carefully tuned by the ECONET consortium in order to reflect the envisaged enhancements to the demonstration prototypes, including the integration of other ECONET solutions not deployed for the final demonstration and the effect of the re-engineering of prototype platforms in more efficient commercial product. More detailed comments can be found in the Annex document.

Also in this case, the overall gain obtained is fully in line with the project targets (77.6% against 80% of energy savings with respect to the BAU scenario), and the energy savings of each network segments appear to span in a range of 54%-83%.

The overall impact in terms of OPEX cut for the considered network scenario is around 303 M€ per year, whose 20% is directly chargeable to the operator infrastructure. This figure suggests the possibility of strongly motivating operators to update their infrastructures to obtain a significant return of investments. A saving of 242 M€ is estimated to spring from the reduction of energy consumption in the end-user homes. In order to obtain these savings, official regulations (like the European Code of Conduct) will have to play a central role in pushing device producers to include advanced green mechanisms in their products for mass markets.

Regarding the environmental aspects, the long term impact of the ECONET project is estimated to reduce the carbon footprint of a medium operator network (including the devices in customer homes) of approximately 1,060 MtCO2e per year, which correspond to the greenhouse gas emissions produced by more than 100,000 cars in a year.

Finally, Fig. 94 reports a summary of the energy savings that can be obtained in the three considered scenarios against the one that was estimated at the beginning of the project activities (included in the DoW and published in [7]), and that was used as reference for setting the project targets (50% of energy saving in the short-term, and 80% in the long-term). As outlined in Section 4.2, this initial estimate was performed by starting from nearly the same reference network infrastructure.

The results in Fig. 94 outlines how the new impact levels derived from the demonstration experimental results are absolutely in line with the initial impact estimates. This is a clear sign of the success of the research and development actions performed in the ECONET project.

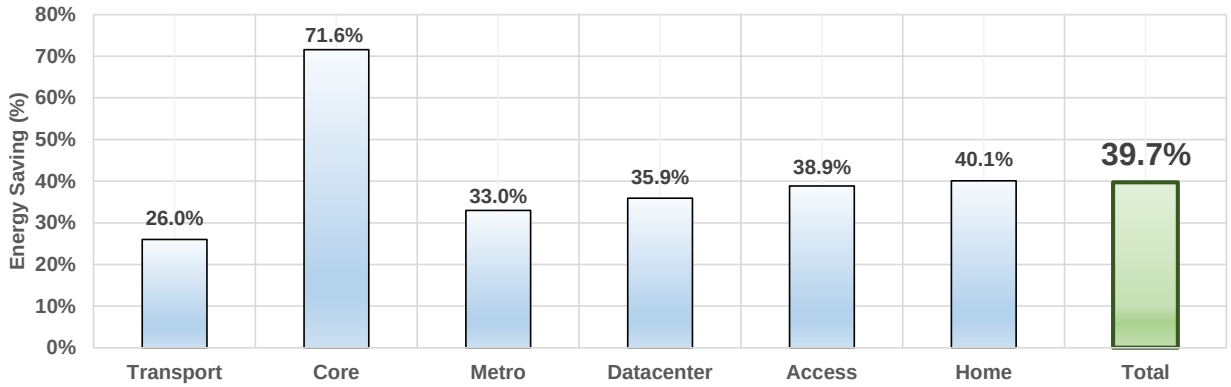

Fig. 85. Estimated energy saving (in %) of the ECONET green demonstration with respect to the BAU scenario, considering the energy profile model defined in Section 4.1 and the experimental results from the demonstration.

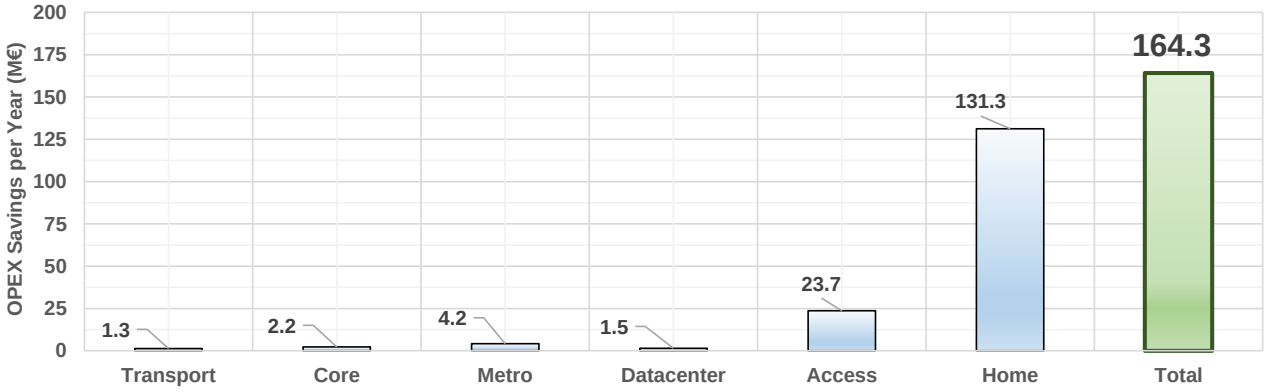

Fig. 86. Estimate of the ECONET green demonstration impact in terms of million Euros savings in operating expenses per year, considering the reference infrastructures defined in [7].

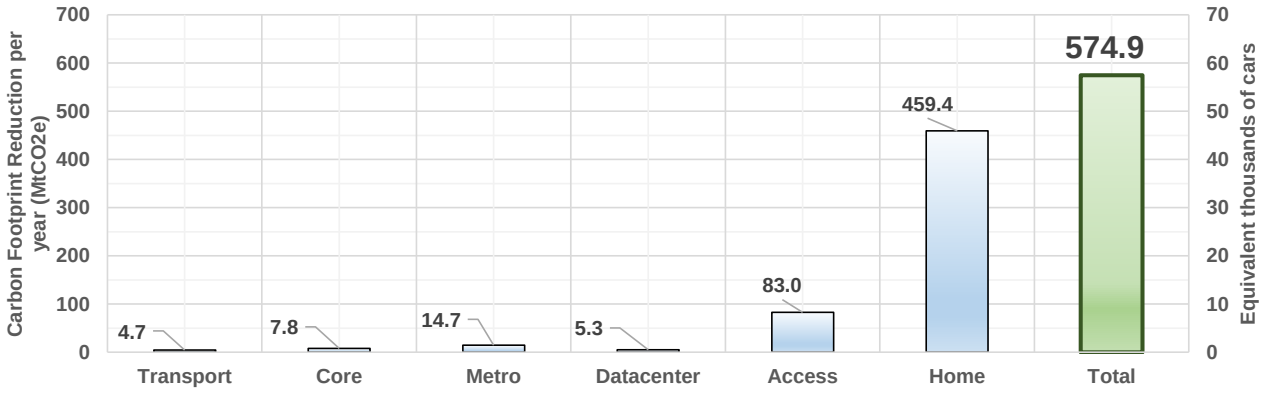

Fig. 87. Estimate of the ECONET green demonstration impact in terms of reduction of the carbon footprint expressed in tons of equivalent dioxite emissions (tCO_2e), and in number of vehicles producing an equivalent amount of emissions [14]. The data labels are expressed in $MtCO_2e$.

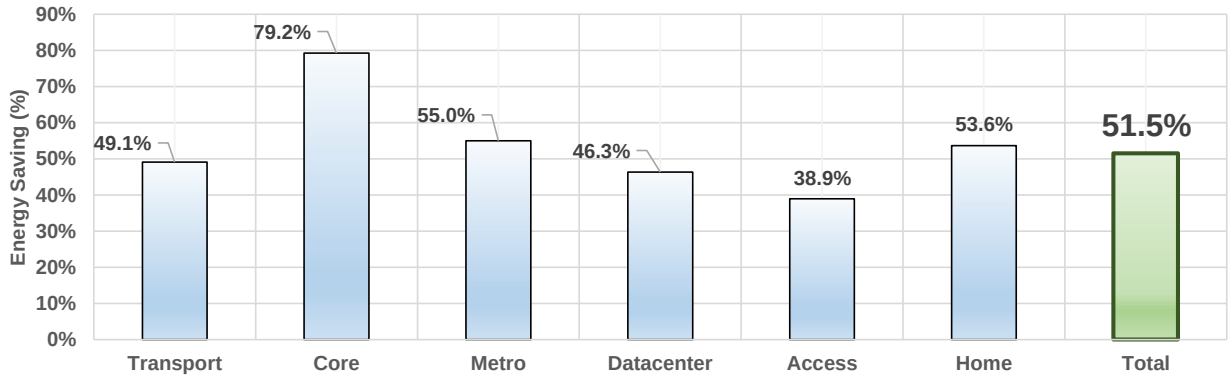

Fig. 88. Estimated energy saving (in %) of the short-term impact of the ECONET tehnologies with respect to the BAU scenario, considering the energy profile model defined in Section 4.1 and the experimental results from the demonstration.

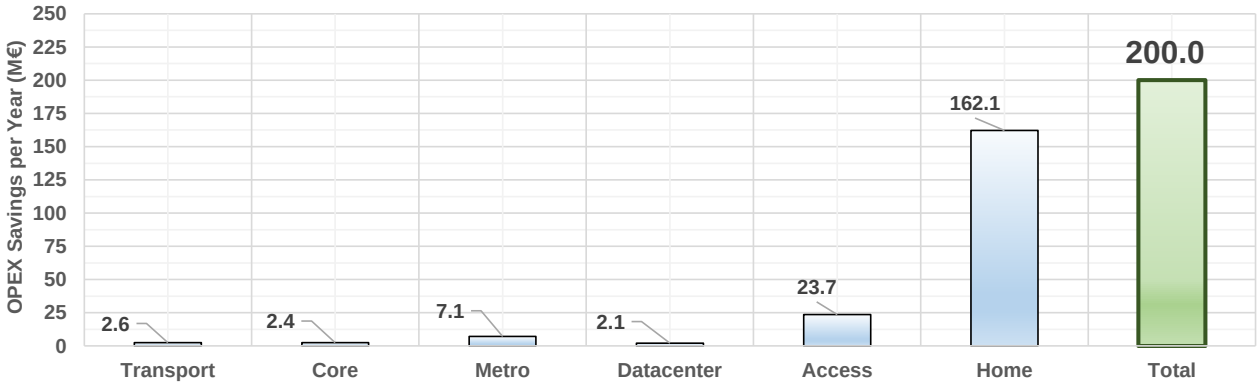

Fig. 89. Estimate of the short-term ECONET impact in terms of million Euros savings in operating expenses per year, considering the reference infrastructures defined in [7].

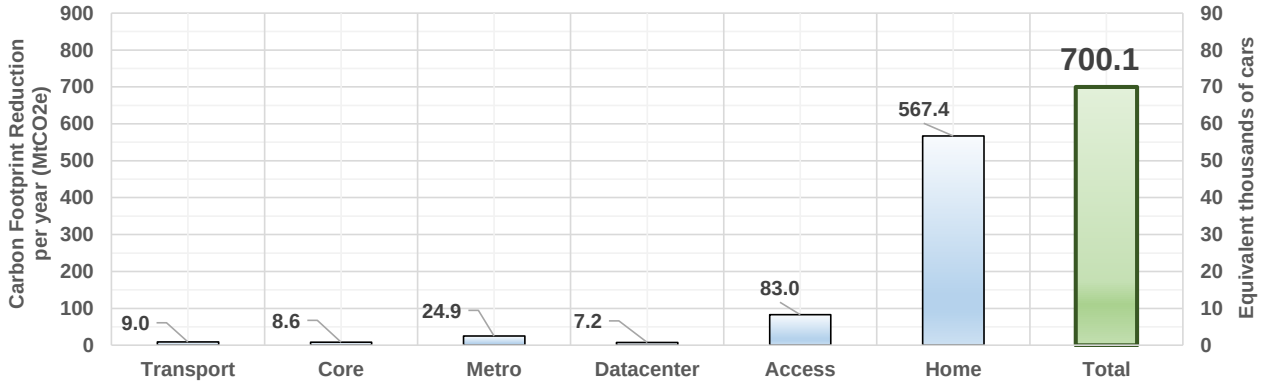

Fig. 90. Estimate of the short-term ECONET impact in terms of reduction of the carbon footprint expressed in tons of equivalent dioxite emissions (tCO_2e), and in number of vehicles producing an equivalent amount of emissions [14]. The data labels are expressed in $MtCO_2e$.

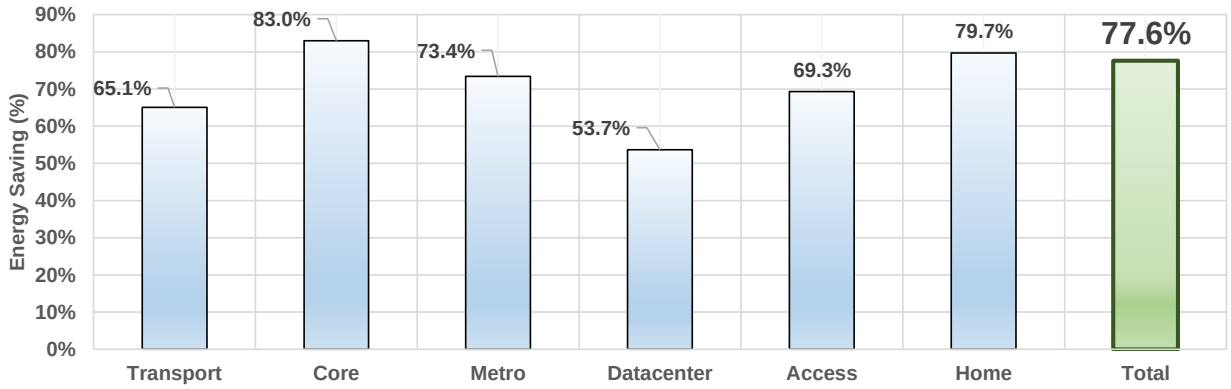

Fig. 91. Estimated energy saving (in %) of the long-term impact of the ECONET tehnologies with respect to the BAU scenario, considering the energy profile model defined in Section 4.1 and the experimental results from the demonstration.

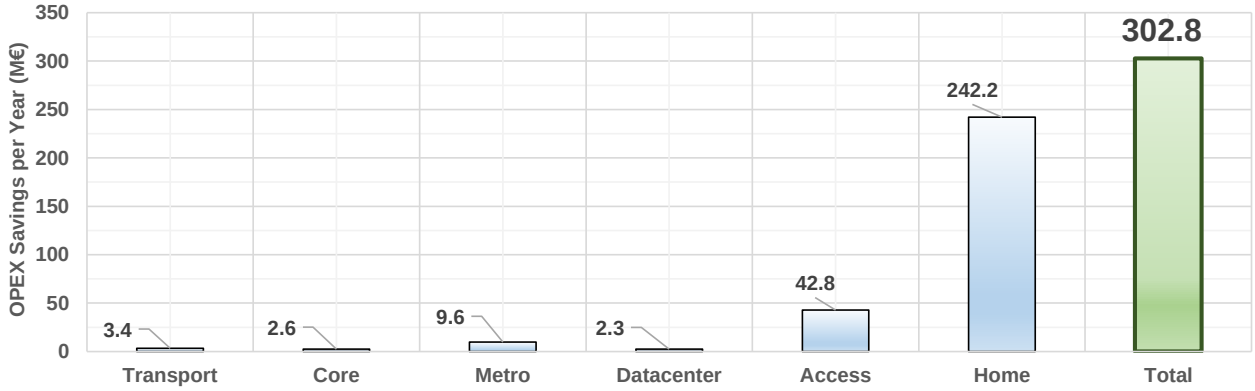

Fig. 92. Estimate of the long-term ECONET impact in terms of million Euros savings in operating expenses per year, considering the reference infrastructures defined in [7].

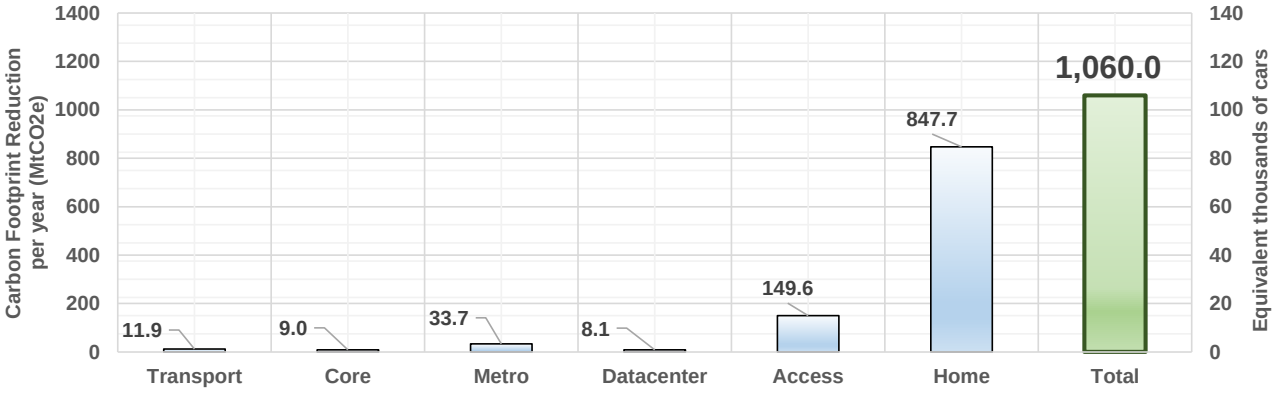

Fig. 93 Estimate of the long-term ECONET impact in terms of reduction of the carbon footprint expressed in tons of equivalent dioxite emissions (tCO_2e), and in number of vehicles producing an equivalent amount of emissions [14]. The data labels are expressed in $MtCO_2e$.

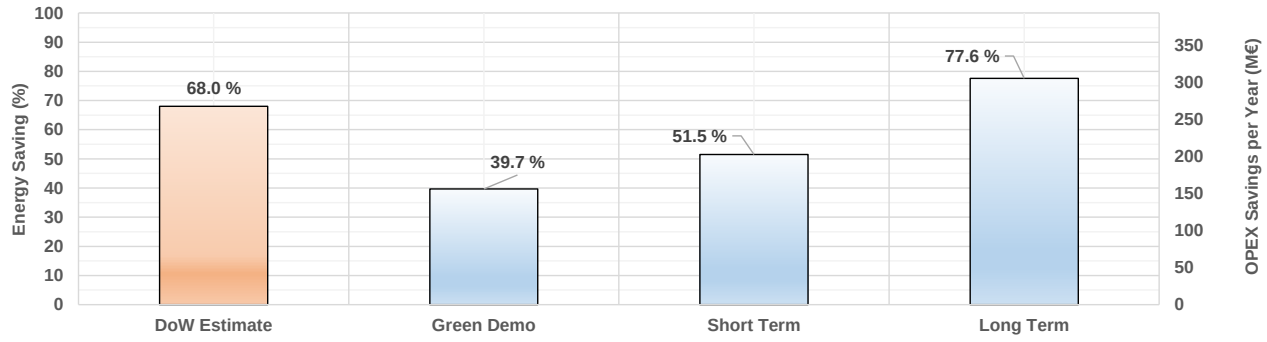

Fig. 94. Summary of the estimated energy and OPEX saving (in % and in M€ per year, respectively) obtained in the three considered scenarios vs the impact estimate performed at the beginning of the project, which was included in the Description of the Work (DoW) and published in [7].

5 Conclusions

This document reported and analysed the experimental results obtained in the final demonstration of the ECONET project, which took place in November-December 2013 in the Telecom Italia site of Largo Borgaro in Turin, Italy.

As described in the D6.4 report, the demonstrator has been specifically designed for representing an entire wire-line network infrastructure of a Telecommunication Operator (Telco) or Internet service Provider (ISP), as well as the homes of end-customers and datacenter networks. For this reason, the demonstrator included 15 physical prototypes of green network devices, in addition to ad-hoc network emulators realized by the ECONET consortium for evaluating the behaviour of green control plane protocols and algorithms.

All the prototypes have been equipped with the most promising green data-plane technologies and solutions proposed in the project WP3. Then, all these data-plane technologies have been exposed through the GAL, to "Local Optimization Policies" and to the ECONET WP5 control and management planes. On these planes, a number of new green protocols (or green extensions to existing protocols) have been used to interface (*i*) energy-aware mechanisms to optimize the trade-off between performance and energy consumption on the entire network, and (*ii*) to expose green metrics and to make them controllable by network operators by means of extended network management and monitoring frameworks. In this respect, it is worth noting that all the demonstrator clouds exposed these metrics by means of the GAL REST protocol, created by the ECONET consortium and defined in the D5.5 report.

During the final demonstration, a number of tests have been performed by emulating typical daily traffic fluctuations, letting the green mechanisms react to the different conditions and traffic levels. Performance metrics have been measured in order to evaluate the energy savings of the green mechanisms, integrated into the prototypes, and the impact on the QoS (in terms of throughput and packet losses) and QoE perceived by end-users (in terms of TCP download times).

Raw experimental results (i.e., as they have been collected during demonstration tests) have been reported and discussed both at the whole demonstrator level and at the level of each cloud. The former aimed at giving a general feeling of what happens to the entire network chain according to the traffic fluctuations. The latter gives fine-grained details on the behaviour and the performance of each integrated green mechanism.

Specific tests have been carried out also to validate the functional behaviour and the performance of the Network Connectivity Proxy. It has been demonstrated that the current version of the Network Connectivity Proxy can effectively and fully replace the network operations of sleeping hosts with a very low overhead in terms of end-user QoE, and of additional power consumption of devices supporting the Proxy.

Then, starting from the raw experimental results, an estimate of the impact of ECONET technologies, as they were set up in the demonstrator, when applied to a Telco infrastructure under realistic operating conditions, has been provided. This estimate also provides a picture on the long-term impact of the ECONET technologies, when the prototypes will be suitably refined to become commercial products.

The results of this estimate highlight how the ambitious goal of reducing the energy consumption of wire-line network infrastructures by 50%-80% has been successfully achieved. In more detail, the impact of the ECONET technologies has been evaluated for the usual operating conditions of a medium-scale Telecom Operator (which corresponds to the 2020 evolution of the network infrastructure of Telecom Italia) under three different scenarios: the scenario corresponding to the final demonstration, the short-term impact scenario, and the long-term one. The obtained results suggest that energy savings with respect to the Business-as-Usual of approximately 39.7%, 51.5% and 77.6%, respectively, can be achieved in the aforementioned cases. These savings can be quantified in a cut of 164, 200, and 302 M€ of the OPEX, and in a reduction of the carbon footprint equivalent to the ones produced by 58, 70 and 105 thousand cars.

References

[1] The ECONET project. “Large-scale validation and benchmarking of network of power-conservative systems”, Deliverable 6.4, https://www.econet-project.eu/.

[2] R. Bolla, R. Bruschi, O. M. Jaramillo Ortiz, P. Lago, “The Energy Consumption of TCP,” Proc. of 3rd ACM/IEEE Internat. Conf. on Future Energy Systems (e-Energy 2013), Berkeley, CA, USA, May 2013.

[3] The ECONET project. “Final integration of device prototypes with energy aware capabilities and local optimization”, Deliverable 6.2, https://www.econet-project.eu/.

[4] R. Bolla, R. Bruschi, A. Cianfrani, M. Listanti, "Enabling Backbone Networks to Sleep ," IEEE Network (IEEE NI), vol. 25, no. 2, pp. 26-31, Mar.-Apr. 2011.

[5] The ECONET project. Network Connection Proxy, WP3 Internal Report no. 7, v. 1.0, https://www.econet-project.eu/Repository/DownloadFile/596.

[6] The ECONET project, Deliverable 3.3, “Final design of green technology for network device data plane”, available on line https://www.econet-project.eu.

[7] R. Bolla, R. Bruschi, F. Cucchietti, K. Christensen, F. Davoli, S. Singh, “The Potential Impact of Green Technologies in Next-Generation Wireline Networks: Is There Room for Energy Saving Optimization?,” IEEE Communications Magazine, vol. 49, no. 8, pp. 80-86, Aug. 2011.

[8] R. Bolla, R Bruschi, A, Carrega, F. Davoli, D. Suino, C. Vassilakis, and A. Zafeiropoulos, “Cutting the energy bills of Internet Service Providers and telecoms through power management: An impact analysis,” Computer Networks, vol. 56, no. 10, pp. 2320-2342, Jul. 2012.

[9] The ECONET project, Deliverable 2.2, “Test plant evaluation criteria and representative test case”, available on line https://www.econet-project.eu/Repository/DownloadFile/312.

[10] The ECONET project, Deliverable 4.1, “Definition of energy-aware states”, available on line https://www.econet-project.eu/Repository/DownloadFile/331.

[11] R. Bolla, R. Bruschi, P. Lago, “The Hidden Cost of Network Low Power Idle,” Proc. of the 2013 IEEE International Conference on Communications (IEEE ICC 2013), Budapest, Hungary, June 2013.

[12] J. Restrepo, C. Gruber, C. Machoca, “Energy Profile Aware Routing,” Proc. IEEE GreenComm’09, Dresden, Germany, June 2009.

[13] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, “An Analytical Model for Designing and Controlling New-Generation Green Devices”, Proc. 3rd Internat. Workshop on Green Communications (GreenCom10), co-located with Globecom 2010, Miami, FL, USA, Dec. 2010.

[14] myclimate – The Climate Protection Partnership, “Offset your car emissions!,” https://co2.myclimate.org/en/car_calculators.