Entity Title Architecture Pilot: Deploying a Clean Slate SDN Based Network at a Telecom Operator

Luiz Cláudio Theodoro,
Pedro Henrique A. D. de Melo,
Flávio de Oliveira Silva,
João Henrique de Souza Pereira,
Pedro Frosi Rosa and Alexandre Cardoso

Federal University of Uberlândia
Uberlândia, MG, Brazil
Email: lclaudio@feelt.ufu.br,
pedrodamaso@mestrado.ufu.br, flavio@ufu.br,
joaohs@ufu.br, pfrosi@ufu.br,
and alexandre@ufu.br

Alex Vaz Mendes,
Murilo Borges Gomes Machado,
Hélvio Pereira de Freitas, Luiz Cláudio Theodoro
and João Henrique de Souza Pereira

Innovation, Research and Development
Algar Telecom
Uberlândia, MG, Brazil
Email: alexvaz@algartelecom.com.br,murilo.bgm@algartelecom.com.br,
helvio@algartelecom.com.br, lclaudio@algartelecom.com.br
and joaohs@algartelecom.com.br

***Abstract*—Due to Internet's remarkable success, its architecture is being challenged to attend new applications requirements such as mobility, security and Quality of Experience (QoE), among others. The requirements that the architecture faces today are far away from the design principles of the protocols in the sixties. Several research initiatives are based on a clean-slate approach, a disruptive view that might result in a completely new network. Our research group proposed the Entity Title Architecture (ETArch), a clean-slate Software Defined Networking (SDN) based approach that currently is able to satisfy applications requirements such as support to multicast traffic, mobility and QoE. This work goes further and presents the deployment of ETArch on a telecommunications service provider network. This work contributes to Future Internet initiatives by presenting a viable approach to deploy new network architectures on top of current providers networks.**

Keywords–Software-Defined Networking;Domain Title Service; Workspace; Telecommunications.

I. INTRODUCTION

Despite the huge success, the Internet architecture is facing a completely new technological context that defies its evolution. In spite of its ubiquity, the Internet has some difficulties to attend new applications requirements, such as mobility, security and Quality of Experience (QoE), among others [1]. The developed protocols contributed for the current Internet success but the requirements that the architecture faces today are far away from their design principles [2] of the sixties.

Several research initiatives [3][4] are on their way in order to provide a solution for the new demands regarding the Internet architecture. One of the approaches to evolve Internet architecture is based on a clean-slate view [5], which suggests drastic changes and might result in a completely new network.

In order to experiment with these new network architectures, several infrastructures are being deployed around the world, such as OFELIA [6] in the Europe, GENI [7] in the United States and FIBRE [8] in Brazil in a joint effort with Europe. These infrastructures enable the deployment and the scaling of the experiments that are necessary to face current Internet scale, however ongoing efforts are using infrastructures that are apart from the current Internet.

Although there are several contributions at a global level, it is not easy to reproduce the research outside a laboratory environment. Experiments involving small equipment sets with few users are an important step for a research validation. However, before going to production, this research must be deployed into real infrastructures and the validation must take into account around millions of users. In this scenario, it becomes a critical issue because the companies do not release their plant and environment in order to be manipulated by researchers with the fear of risking themselves because somehow could affect the services provided for the public.

In previous work, our research group proposed the Entity Title Architecture (ETArch) [9], a clean-slate Software Defined Networking (SDN) based approach which aims at satisfying different applications requirements, such as support to multicast traffic [10] , mobility [11], and QoE [12].

The present work goes further and presents and details the deployment of ETArch at a telecommunication service provider production network. This initiative represents an important path towards the actual deployment of Future Internet initiatives on real networks.

ETArch architecture guarantees the possibility of implementing a clean slate SDN network, in a telecom operator, by using the concept of horizontal addressing based on titles, an unambiguous designation of an entity. In a traditional scenario, involving residential customers connected to a commercial network, an alteration in their access methods is induced in a non-conventional way. All of this happens without great setbacks for the final user and a ETArch based chat application is executed with traffic monitoring among the participants.

After this deployment, there is a firm intention to promote a set of developments which can attest the efficiency and potentiality of this solution collaborating so that the future Internet can answer to a series of current and future demands.

The remainder of this work is organized as follows: Section II presents an overview of related work about deployment of new network architectures on top of current networks. Section III introduces ETArch concepts and presents the operator infrastructure. Section IV describes the technical aspects of the deployment. Section V presents some results of the current work and finally, Section VI presents some concluding remarks and future work.

II. Related Work

Several research groups has initiatives in order to make Internet capable to support the new requirements that challenges the current Internet architecture. One of the approaches is to decouple the architecture and the infrastructure and the OpenFlow standard is cited as one of the most popular solutions to this end [13][14][15]. Some of these initiatives goes towards the direction to deploy the research results onto real environments.

BeHop [16] has an interesting approach, by implementing a wireless testbed for dense WiFi networks frequently found in residential and corporative environments. This prototype was implemented in a campus with about thirty active users as "guinea pigs" allowing researchers to study and evaluate pros and cons of new ways of controlling WiFi networks. BeHop supplies a general purpose framework for experimenting new techniques in order to control power, channel allocations and associations.

The integration between BeHop and the production network showed the benefits of an implementation in the actual world while maintaining the aimed flexibility to process others experiments. It was essential to study and to evaluate the WiFi management strategies and its impacts on the conditions found in a real network such as clients diversity, mobility, and interference with neighbor networks. A testbed was used in order to transport real traffic of users connected to WiFi devices and at the same time to keep the flexibility to apply frequent changes and occasionally force the network down aiming to show network resilience.

Yiakoumis' work [17] points out the possibility of leaving the network control to the user instead of the Internet Service Provider (ISP). This statement raised controversies and contrary opinions have been coming up but the defended idea is to allow that the user's choice can guide the network traffic managing not only inside residences but also inside the ISP [18].

An interesting implementation has been made by Hampel, Steiner and Bu [19]. They suggest the idea of the SDN in an operator, but on top of an Internet Protocol (IP) network. In this case, OpenFlow capable elements run vertical forwarding to interoperate with a legacy infrastructure using IP in consonance with routing traditional protocols.

All these proposals have in common the approach that new solutions could be created by using current network infrastructure. These new solutions are decoupled from the network infrastructure and enables new types of experiments using a SDN based approach.

This study also is based on the assumption that SDN is the enabler of changes in the network which would make it more programmable and flexible. However, it goes further and deploys a clean-slate network architecture on top of a legacy infrastructure at a real network operator. A particular feature of ETArch is that it aims at supporting the application communication requirements over time and support these requirements from top to lower layers of the protocol stack.

III. Entity Title Architecture (ETArch) Pilot

Countless researches are being made with the intention of recreate the Internet architecture that collaborates for the evolution of this great worldwide network. The more expressive proposals have built a large-scale experimentation facility, supporting both research on networks and services, by gradually federating existing and new testbeds for emerging or future Internet technologies [20][4]. Joining this researchers initiative around the world, the ETArch Pilot group intends to create conditions so that researches in future experiments leave the laboratory and go to the actual world.

The suggested architecture in this work has as basic point of view to semantically approximate upper and lower layers and for that it uses the ETArch proposal operating over a commercial network. When it comes to this model it foresees the possibility of attending Internet demands whether current or future ones. ETArch is a clean slate network architecture in which identification and addressing schemes are based in an independent topology designation that uniquely identifies an Entity: its Title. ETArch transport mechanism is based on a logical channels named Workspace which is capable of unifying multiples communication entities.

Entities in the Title Model [21] differ from the defined concept in some literatures and they are not considered simple resources inside a network but beings whose communication needs must be understood and supported by the Service Layer and then by the lower layers as Physical and Link layers. Hereinafter ETArch main concepts and components are better detailed.

An Entity has list of requirements and capacities related to communication. An entity may be an equipment, a user, an application, a thing, and so on. It has at least one title and one location, known as Point of Attachment (PoA). From mobility point of view such separation is important because an entity's location could change over time. These entities can relate among themselves and through such relations they can inherit properties, except the title [9][22].

A group of Titles is bundled in a Namespace which also must have a single title. A Title can be represented by a tuple and its specification could be such as Namespace::Identification-entity. On the other hand, Workspace is a logic bus, independently of underlying topology, with which entities can be attached to be part of a communication domain. The entities addressing happens in an application level and these entities do not communicate directly but they communicate through a workspace.

The Workspace can work on wired and wireless networks. The workspace has the following properties: a title; a group of Network Elements (NE); a list of capacities; the visibility; and a list of requirements. The title identifies that workspace in a unique way. The NE list represents the physical infrastructure that supports the workspace. The capacities indicate the properties that the workspace must satisfy regarding communication such as QoS and security parameters. The requirements must be supported by one entity that wants to attach to a given workspace.

A Workspace is created when an entity produces some kind of thing that can be consumed by other entities such as in a file sharing, a content or a video streaming. It is controlled by a Domain Title Service (DTS), which has the responsibilities

 ISBN: 978-1-61208-411-4

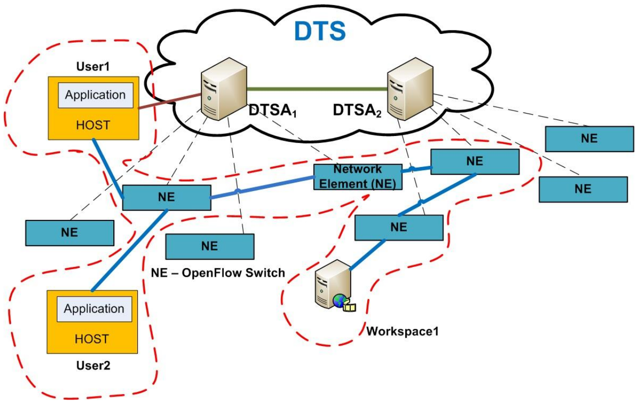

Figure 1. Components of the Architecture - DTS, DTSA, Entity, Title and Workspace.

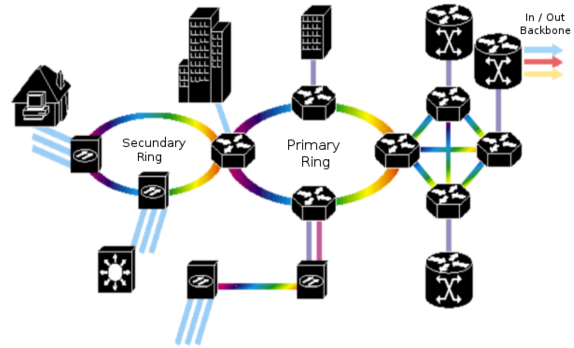

Figure 2. MetroEthernet Topology

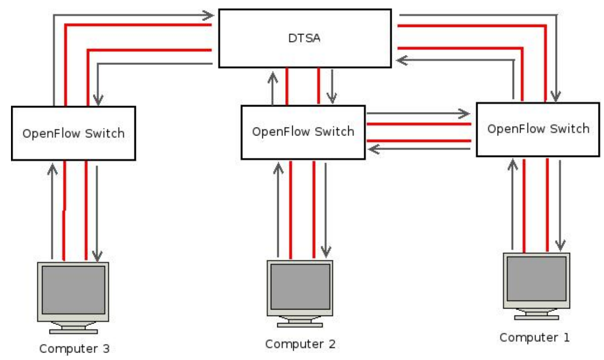

Figure 3. OpenFlow Network

mentioned as follows: titles resolution, entities management and their relationship with their Entities and Titles relation management. The DTS is constituted by DTS Agents (DTSAs) and it has a base of knowledge with information from the environment so it is possible to monitoring and controlling the requisites from an entity. The DTSA manages the entities cycle of life from the beginning to the activities end. The Figure 1 shows an illustration of a Workspace controlled by DTSAs encompassing some NEs.

IV. Deployment and Experimentation

This work proposes the deployment of a testbed based on ETArch architecture as a response to the challenges regarding future internet. By considering that current Internet is an ubiquitous architecture, it is desirable that any changes would be transparent to the users around the world.

By having that ETArch can be deployed on openflow switches, it is possible to implement it by using a telecom operator MetroEthernet network. At the last mile, OpenFlow based switches will be deployed as customer premise equipments.

The operator where this work was conducted is Algar Telecom [23]. A Brazilian operator located in the southeast region of the country. Considering 2014 information, the company has 1.321 million customers and 380 thousand customers in the broadband access. The deployment in this work considers the technologies related to this group of broadband access users.

The MetroEthernet network topology is based on primary and secondary rings as depicted in the Figure 2. Home users are connected to the secondary rings. The primary rings are distributed over the interconnected area and link the secondary rings. Such equipments in their vast majority operate with Synchronous Digital Hierarchy (SDH) or Dense Wavelength Division Multiplexing (DWDM) [24]. Primary rings are characterized by throughput superior to 100 Gbps. Secondary rings link smaller geographic area and their throughput are under 40 Gbps.

To provide access to the customers in the last mile, the operator uses different access technologies such as Asymmetric Digital Subscriber Line (ADSL), Very-high-bit-rate Digital Subscriber Line (VDSL) and Fiber To The Home (FTTH) [25]. Usually the customers are connected to the secondary rings. In some cases, the customers requires higher throughput rates such as in a private condominium, a commercial building or a big company. These customers can be connected to the primary ring as depicted in the Figure 2. In typical ADSL access, the modem is connected to a Digital Subscriber Access Method (DSLAM), which in turn is attached to the secondary ring in a Metro Network. Secondary rings are aggregated in primary ring that connects to the IP routing core.

Two or more areas (domains) will be linked by a virtual networking, by using IP networks, providing researches and developers with an actual usage scenario for the deployments and tests. At a local environment, inside an area, OpenFlow switches are directly linked to each other, as represented by Figure 2.

However, as shown in the operator's network the OpenFlow switches would be separated by an IP structure that does not implement such specification conforming Figure 3.

There are several challenges related with the deployment of ETArch on a MetroEthernet network infrastructure. For example, ETArch does not use the TCP/IP protocol stack on its control and data planes. The legacy infrastructure is completely based on TCP/IP. To solve the problem it was necessary to use a strategy where the ETArch components could communicate in a transparent way over the legacy infrastructure. By using this approach, the legacy infrastructure would work only as a forwarding plane. As a result, the OpenFlow switch will work as being directly connected, as depicted in the Figure 3.

The forwarding graph in an OpenFlow based network can be discovered by using the Link Layer Discovery Protocol (LLDP) [26]. Briefly, LLDP is a link layer protocol used

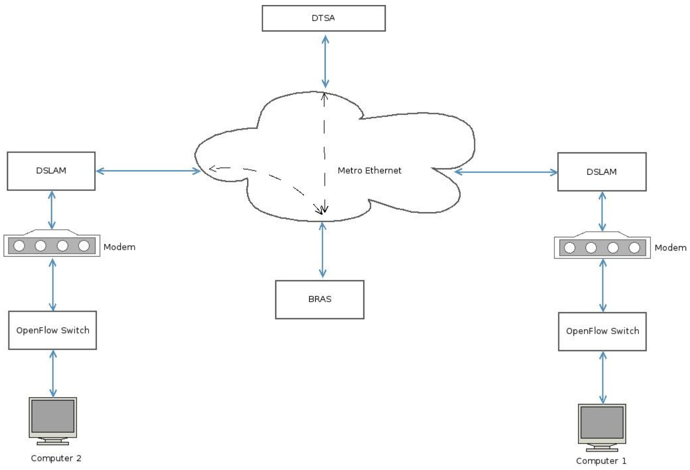

Figure 4. Operator Network

by network devices for informing their identity on a local area network (LAN). The forwarding graph contains all of the OpenFlow switches that are directly connected. However, the legacy infrastructure contains several middleboxes between the switches. These switches are also geographically scattered along the network infrastructure.

Therefore, it was necessary to solve this issue to have all switches directly connected in a virtual way. The adopted solution was to create a tunnel using the Generic Routing Encapsulation (GRE) protocol [27] in order to forward the LLDP frames. The GRE protocol is very popular due to its simplicity and compatibility. GRE is a tunneling protocol that can encapsulate, among other things, the link layer protocols inside a virtual point-to-point link over an IP network. Thus, it was possible to create a virtual network of OpenFlow switches directly connected over the operator IP network, according to Figure 4.

Thus, the tunnel creation could interconnect the OpenFlow switches. However, the telecom operator does not have OpenFlow capable switches deployed on the infrastructure. Replacing current switches is not a feasible alternative. To overcome this, Open vSwitch (OVS) was used. Although this brings a new requirement for the deployment, OVS can be easily installed in commodity hardware at customer premises. OVS is an open-source implementation of a virtual multilayer switch and support multiple protocols and standards including some OpenFlow versions.

In the end, the virtual switches, based on OVS, where scattered along the network and interconnected using GRE tunnels. The DTSA acts as the OpenFlow controller of these switches.

The deployment was based on the most recent version of the software components of ETArch. The components were installed in servers located inside the Operator network. In the first stages, PCs using Ubuntu Linux (14.04) operating systems played the role of each customer's equipment. On each PC, OVS was installed and a new bridge interface (br0) was created. This bridge interface can be used also for data plane and for the control plane in order to communicate with the DTSA.

From the moment this process is realized in both machines, the bridge is added to the switch and established a register in the controller. Both switches (machines) use a GRE because it offers a tunnel, by simulating a link between two network nodes, and this is done by using an IP address (then attending to the initial proposal). Atop of it, the OVS itself is already capable of creating this kind of tunnel internally thus the process is finished by adding the GRE tunnel in the same bridge of the controller.

V. Results

To meet the initially proposed objectives, the verification may be done by the assembly of the network structure and by the execution of an application specifically developed for the ETArch architecture.

The first scan mode is based on the application of ETArch controller which asks every registered switch to send LLDP as a regular package. This proceeding is very important for the network mapping itself because it signals what are the most important options and the paths between them.

In this way, the controller can take a complete view of the network. A packet monitoring tool (tcpdump) has been used to observe the exchange of information between switches, specifically the LLDP messages, encapsulated in GRE packets as mentioned before, as shown in Figure 5.

Such information indicates that these options are virtually adjacent; which means that the legacy structure remains present in the network lab, transparently, where it was initially deployed, which did not require considerable changes in the usual practices within the carrier.

Upon confirmation of adequate controller setup, the switches (machines) were transferred to MetroEthernet environment, where it obtained a valid IP. Every customer has been connected by using the existing network, meaning that customers would be connected to each other by the controller and network elements such as switches, DSLAMs and the access modems. For the construction of GRE tunnels, the OVSs must have a valid IP and therefore the customer must send the Ethernet over the Point-to-Point Protocol (PPPoE) frames to the ISP.

The OVSs setup process is repeated by replacing the previous IPs, by the one obtained on the new network, and thus all the legacy structures are transparent for the control and connection between customers. LLDP packets and their respective answers are inspected to check whether this phase of the process has been successful.

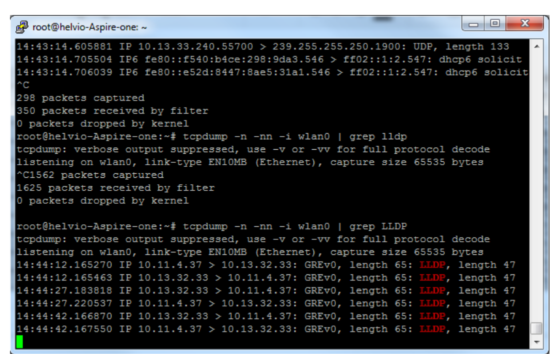

Figure 5. Monitoring LLDP

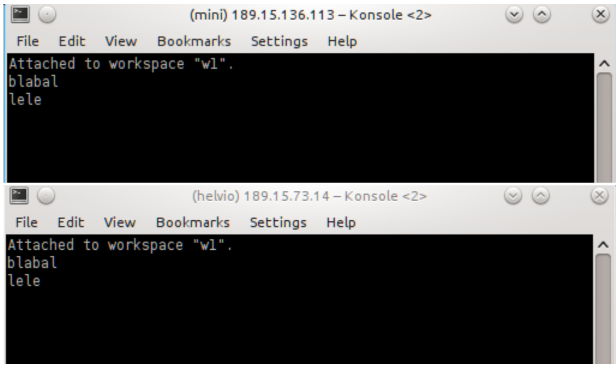

Figure 6. Running Chat

A chat application is invoked to prove that the deployed environment, which involves entities and other ETArch communications concepts[28], is working properly (even though it is using non-traditional TCP/IP protocols).

The chat application has been developed in Python and invoked directly on a command line interface, by passing parameters as entity titles and Workspace. By invoking them (chats) on both machines, each one by its title, it was found that, according to the concept defined in the ETArch architecture, both applications could share the specified workspace (W1), as it can be seen in Figure 6. It allows validate the implementation and shows the results obtained by the new architecture in the proposed environment.

VI. Concluding Remarks and Future Work

This work presented the deployment of ETArch, a clean slate network architecture, over the production network of a telecom operator. In order to test and experiment this deployment, a chat application was used. This application is based on ETArch concepts such as Workspaces, Entities and Titles enabling the use of a new protocol stack between end users over the legacy networks.

In this process, the physical infrastructure was kept in place and only software based framework where added to the infrastructure. In the end user side, a software based OpenFlow switch was introduced and on the operator side, the DTSA, the entity responsible for the control plane of ETArch.

As a future work there are several previewed fronts such as increasing the number of customers and switches, the deployment of new applications based on ETArch workspaces in order to show its efficiency in areas such as video streaming and finally the withdraw of some network elements from the legacy network.

We are confident that this deployment will bring facilities and dynamism to researchers facing the Future Internet's evolution and it can enable new types of services and applications which can be offered by the operator to their customers, helping to bring Future Internet research into reality.

Acknowledgment

This work has been partially funded by the European Community's Seventh Framework Programme, under grant agreement n. 258365 (OFELIA project), by the Brazilian agencies: CAPES, CNPq and FAPEMIG and also by PROPP/UFU. We also would like to thanks ALGAR Telecom for the support and partnership on this work.

References

[1] T. Zahariadis et al., "Towards a future internet architecture," in The Future Internet. Future Internet Assembly 2011: Achievements and Technological Promises, ser. Lecture Notes in Computer Science, J. Domingue, A. Galis, A. Gavras, T. Zahariadis, and D. Lambert, Eds. Berlin, Heidelberg: Springer-Verlag, May 2011, vol. 6656, pp. 7–18. [Online]. Available: http://www.springerlink.com/content/978-3-642-20897-3#section=881237&page=15&locus=86

[2] D. Clark, "The design philosophy of the darpa internet protocols," SIGCOMM Comput. Commun. Rev., vol. 18, no. 4, Aug. 1988, pp. 106–114. [Online]. Available: http://doi.acm.org/10.1145/52325.52336

[3] J. Pan, S. Paul, and R. Jain, "A survey of the research on future internet architectures," Communications Magazine, IEEE, vol. 49, no. 7, July 2011, pp. 26–36.

[4] A. Gavras, A. Karila, S. Fdida, M. May, and M. Potts, "Future internet research and experimentation: the fire initiative," ACM SIGCOMM Computer Communication Review, vol. 37, no. 3, 2007, pp. 89–92.

[5] J. Rexford and C. Dovrolis, "Future internet architecture: clean-slate versus evolutionary research," Communications of the ACM, vol. 53, no. 9, 2010, p. 36–40.

[6] OFELIA. OpenFlow in europe - linking infrastructure and applications. [Online]. Available: http://www.fp7-ofelia.eu/about-ofelia/ [retrieved: May, 2014]

[7] C. Elliott, "GENI: opening up new classes of experiments in global networking [Internet predictions]," IEEE Internet Computing, vol. 14, no. 1, Feb. 2010, pp. 39–42.

[8] FIBRE. FIBRE Project - Future Internet Testbeds Experimentation Between Brazil and Europe. [Online]. Available: http://www.fibre-ict.eu/ [retrieved: May, 2014]

[9] F. de Oliveira Silva, J. H. de Souza Pereira, P. F. Rosa, and S. T. Kofuji, "Enabling future internet architecture research and experimentation by using software defined networking," in Software Defined Networking (EWSDN), 2012 European Workshop on. IEEE, 2012, pp. 73–78.

[10] M. Amaral Gonçalves, F. de Oliveira Silva, d. S. Pereira, J. Henrique, and P. Frosi Rosa, "Multicast Traffic Aggregation through Entity Title Model," Jul. 2014, pp. 175–180. [Online]. Available: http://www.thinkmind.org/index.php?view=article&articleid=aict_2014_7_40_10177

[11] C. Guimaraes et al., "IEEE 802.21-enabled Entity Title Architecture for Handover Optimization," in 2014 IEEE Wireless and Communications and Networking Conference. Piscataway, NJ:: IEEE, 2014, pp. 2671–2676.

[12] Castillo et al., "Evolving Future Internet Clean-Slate Entity Title Architecture with Quality-Oriented Control Plane Extensions," Jul. 2014, pp. 161–167. [Online]. Available: http://www.thinkmind.org/index.php?view=article&articleid=aict_2014_7_20_10164

[13] B. Raghavan et al., "Software-defined Internet architecture: Decoupling architecture from infrastructure," Journal of Something, 2012, p. 43–48.

[14] T. Anderson et al., "NEBULA - A Future Internet That Supports Trustworthy Cloud Computing," 2012, pp. 1–31.

[15] J. Roberts, "The clean-slate approach to future internet design: a survey of research initiatives," annals of telecommunications - annales des télécommunications, vol. 64, no. 5-6, May 2009, pp. 271–276. [Online]. Available: http://dx.doi.org/10.1007/s12243-009-0109-y

[16] Y. Yiakoumis et al., "BeHop: a testbed for dense wifi networks," in Proceedings of the 9th ACM international workshop on Wireless network testbeds, experimental evaluation and characterization. ACM, 2014, pp. 1–8.

[17] ——, "Putting home users in charge of their network," in Proceedings of the 2012 ACM Conference on Ubiquitous Computing. ACM, 2012, pp. 1114–1119.

[18] M. Chetty, R. Banks, A. Brush, J. Donner, and R. Grinter, "You're capped: understanding the effects of bandwidth caps on broadband use in the home," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, 2012, pp. 3021–3030.

[19] G. Hampel, M. Steiner, and T. Bu, “Applying software-defined networking to the telecom domain,” in Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on. IEEE, 2013, pp. 133–138.

[20] J. Rexford and C. Dovrolis, “Future internet architecture: Clean-slate versus evolutionary research,” Commun. ACM, vol. 53, no. 9, Sep. 2010, pp. 36–40. [Online]. Available: http://doi.acm.org/10.1145/1810891.1810906

[21] J. de Souza Pereira, F. de Oliveira Silva, E. Filho, S. Kofuji, and P. Rosa, “Title model ontology for future internet networks,” in The Future Internet, ser. Lecture Notes in Computer Science, J. Domingue et al., Eds. Springer Berlin Heidelberg, 2011, vol. 6656, pp. 103–114. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-20898-0_8

[22] F. d. O. Silva, J. H. d. S. Pereira, S. T. Kofuji, and P. F. Rosa, “Domain title service for future internet networks,” in Anais do II Workshop de Pesquisa Experimental na Internet do Futuro (WPEIF). Campo Grande: SBC, 2011, pp. 33–36.

[23] TELECO. Algar telecom. [Online]. Available: http://www.teleco.com.br/en/en_operadoras/ctbc.asp [retrieved: May, 2015]

[24] M. Huynh and P. Mohapatra, “Metropolitan ethernet network: A move from lan to man,” Computer Networks, vol. 51, no. 17, 2007, pp. 4867–4894.

[25] A. S. Tanenbaum, “Computer networks 4th edition,” ed: Prentice Hall, 2003.

[26] P. Congdon, “Link layer discovery protocol and mib,” V1. 0 May 20. 2002, http://www. IEEE802, 2002.

[27] S. Hanks, D. Meyer, D. Farinacci, T. Li, and P. Traina, “RFC 2784 - generic routing encapsulation (GRE),” 2000.

[28] F. de Oliveira Silva et al., “Semantically enriched services to understand the need of entities,” in The Future Internet, ser. Lecture Notes in Computer Science, F. Álvarez et al., Eds. Springer Berlin Heidelberg, 2012, vol. 7281, pp. 142–153. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-30241-1_13