Virtual CPE: Enhancing CPE's deployment and operations through Virtualization

Pau Minoves
pau.minoves@i2cat.net
i2CAT Foundation

Ole Frendved
ole.frendved.hansen@uni-c.dk
UNI-C

Bo Peng
bpeng@essex.ac.uk
University of Essex

Andrew Mackarel
andrew.mackarel@heanet.ie
Dave Wilson
dave.wilson@heanet.ie
HEAnet

***Abstract*— The arising opportunity exposed in this paper is leveraging the capabilities of network virtualization to provide a whole IP Network as a Service (NaaS). Users are offered the capability to customize the policies, the configuration and the management on their allocated virtual network to meet their own requirements. Following the NaaS paradigm, the European Community funded project Mantychore is developing and providing an IP NaaS framework to enhance the service portfolio of National Research and Education Networks (NRENs) and other e-Infrastructure providers. This paper, as a particular use case of the project, describes how NaaS could benefit Customer Premises Equipment (CPE) deployment. Two CPE scenarios discussed are: the Aggregated Virtual CPE and the Multiple Networks Virtual CPE. The benefits of these scenarios go from equipment purchasing cost to operations expenses and efficiencies.**

***Index Terms*— Network as a Service, network virtualization, CPE.**

I. Introduction

A recent challenge for network operators is the deployment of dynamic network infrastructures capable of supporting many different types of application, each with their own access and network resource usage patterns. The network as a Service (NaaS) paradigm is introduced recently as a key enabler to address this challenge by offering virtual infrastructures to third party through software and/or web service. Network virtualization plays a crucial role in this paradigm. It is the creation of a virtual version of a physical resource (e.g. network, router, switch, optical device or computing server), based on an abstract model of that which is often achieved by partitioning (aka slicing) and/or aggregation.

On traditional hardware-based CPE, deploying fiber-based connectivity to an institution requires several layers of equipment to be deployed, much of it on site at the institution. Some of this equipment is needed in order to provide a network boundary between the provider and the client institution (Client Premises Equipment: CPE). The goal of the Virtual CPE use case is to rationalize that equipment, by investigating whether a single device at the client institution can share the role, or if it can be provided by a small number of aggregating devices at a central location.

European FP7 Mantychore project [1] follows the NaaS paradigm to enable National Research and Education Networks (NRENs) and other e-Infrastructure providers to enhance their service portfolio by building and deploying software and tools to provide infrastructure resources like routers, switches, optical devices, and IP networks to virtual research communities. It provides an architectural solution for virtual IP network for network infrastructure provisioning. Along with these new service scenarios, a comprehensive toolset that implements these visions into actionable software, named OpenNaaS[7], is under development at the Mantychore project. OpenNaaS is a modern and flexible NaaS middleware that empowers service providers with increased and sophisticated management flexibility on the network.

This paper describes how the Network as a Service approach being taken by the Mantychore FP7 project can use network virtualization technologies to enhance current Customer Premises Equipment (CPE from now on) installations. In order to do so, two CPE use cases and their enhanced Virtual CPE equivalent will be exposed (or vCPE for short). In the first one, Virtual CPE for educational users, emphasis is placed on defining the demarcation between provider and user; in the second, Multiple Network Support, emphasis is placed on the defining the demarcation between two providers towards a single user.

The rest of the paper is organized as follows. Section II introduces our proposal in the context of the research undertaking in the Mantychore project and provides a brief overview of the related research projects in network virtualization area. The following section, Technical Approach, further explores the technological implementation of the Network as a Service concepts exposed and how the Virtual CPE ideas are instantiated. The section after, Enhancing CPE, will describe the detail of the scenarios under study as well as a comparative discussion with the proposed Virtual CPE equivalent. Furthermore, in Future Work an exposure of the medium term roadmap and planned deployments of the Virtual CPE scenarios is done.

II. Related Work

One of the main benefits of network virtualization is the separation of roles that it enables. Typically, three roles can be identified in the NaaS scenario, according to 4WARD[2]:

- **Infrastructure Provider (InP):** the infrastructure owner who deploys and maintains the network equipment.
- **Service Provider (SP):** harvests network instances from one or more InPs and integrate them into a management domain. SP provides various services to end users (e.g. an IP Network service).
- **End User:** uses the services offered by SP.

InPs are able to partition their physical infrastructures by adopting NaaS framework, based on user/application requirements into virtual networks, and offer them as network services to service providers. A virtual infrastructure is a set of virtual resources interconnected together and managed by a single administrative entity. The logically independent virtualized infrastructure provides a flexible pay as you go method that allows SP manage their resources efficiently on the dynamic demand from various end users.

There are a number of initiatives both in Europe and around the world, which are related to network virtualization. The European FP7 4WARD [3] project is a wide-scope project, aiming to develop clean-slate solutions to the current ossification of Internet in a number of different areas. 4WARD developed a systematic and general approach to network virtualization framework that allows the coexistence, inter-operability, and complementarity of several network architectures both in fixed and wireless network. In 4WARD framework, scalability was a major challenge that is considered, particularly in terms of provisioning, management and control of virtual networks. To achieve this, a number of essential algorithms for scalable mapping and embedding of virtual resources into the infrastructure, including discovery, matching, and binding were developed [10-13]. Performance evaluation suggested that the provisioning of virtual networks at large scale is, in principle, feasible. Compared with 4WARD, Mantychore is designing and implementing a mechanism for network virtualization and provisioning, while also proposing a novel resource brokering mechanism for fast and efficient resource trading among multiple InPs. In addition, Mantychore provides a functional and pre-operational deployed software product to enhance network operator and SP‘s service portfolio, while 4WARD aimed at a proof of concept.

GEYSERS [4] is another European FP7 project which proposes a novel architecture to support 'Optical Network + IT' resource provisioning for end-to-end service delivery. GERSERS develops a mechanism that allows infrastructure providers to partition their resources (optical network and/or IT), composes specific virtual infrastructures and offers them as a service to network operators. An enhanced network control plane capable of provisioning advanced optical network service as well as IT resources controls each virtual infrastructure. This service delivery is realized in a four-layer architecture that aims at opening new business models not only to service providers but also allow infrastructure providers to offer infrastructure services on demand. This is achieved through a Logical Infrastructure Composition Layer (LICL) capable of abstracting, partitioning and composing virtual infrastructures from a set of physical resources in an automated way. Similar to 4WARD project on the vision of network convergence and unification, GEYSERS is more application-driven, and target to provision virtualized IT and optical network resources to support on-demand applications. Compared with GEYSERS, Mantychore service does not restricted to any particular layer in the network architecture, and therefore, provides a comprehensive cross-layer virtualization solution for NaaS provisioning. In this Virtual CPE use case, Mantychore FP7 aims to building functional IP Networks.

The FP7 FEDERICA project [5] is a sliceable e-infrastructure dedicated to European network researchers. FEDERICA uses resources (e.g. routers, switches, computing nodes) contributed by its partners and interconnected through GEANT [6] links as the supporting substrate. When a research user wants to perform an experiment, the Network Operations Centre (NOC) allocates a slice of infrastructure resources to him (e.g. computing nodes, physical or software routers), and let the user control it to perform his own experiments. Unlike FEDERICA, Mantychore does not deploy a new e-infrastructure; it allows the existing NREN e-infrastructures to provide a new service to the research community

The US PlanetLab project [8] is a system for managing and controlling network infrastructure that accommodates experiments running for short period of time in parallel with live services. PlanetLab provides isolation of services by running each service using virtualized set of network resources. One important goal of PlanetLab is to increase the number of users who can work on a single node. The PlanetLab users are allowed to create their own slices, which is a set of virtual machines distributed over PlanetLab nodes. Notably, PlanetLab does not provide full isolation of resources (e.g., routers). The increase in the number of virtual networks degrades the performance. Bandwidth reservation per virtual infrastructure is not feasible in PlanetLab, which leads to the contention of link bandwidth.

AKARI [9] is a Japanese project for designing next generation network architecture, which is supported by the National Institute of Information and Communications Technology (NICT) of Japan. Network virtualization has been considered as a major enabler for their network design and architecture. In particular, they proposed a virtualization layer where the problem of how to isolate and allocate existing physical resources has been considered. Additional requirements regarding virtualization include the flexibility of abstraction layer, mapping using multiplexing technology, resource management and interconnections. However, at the time of writing this paper, only limited information is publicly available from the AKARI project on the technical details of how the virtualization will be implemented.

III. Technical Approach

In order to achieve an operative deployment of the Virtual CPE so that network operators can exploit its benefits an appropriate toolset needs to be built. The Mantychore FP7 project main development effort is focused into building the OpenNaaS framework. This is a custom NaaS-oriented framework that contains the technological foundation to implement ambitious NaaS applications. The Virtual CPE application will thus be build on top of this framework, as an extension.

This framework is, however, build around versatile concepts so several NaaS applications can reuse it without having to accommodate to the same strict NaaS orientation. The OpenNaaS framework is build around very simple concepts that can accommodate several orientations. They can be summarized as follows:

- **Resources**: The basic unit of functionality in OpenNaaS is the resource. This can be anything that a user can own and operate. Both a physical and a logical router can be resources, as well as a BoD domain, an IP network or ROADM equipment.
- **Model**: All resources contain a model. This allows inspecting the internal state of the resource, as well as its capabilities and functionalities. Changes on the resource state are reflected on the model. This is the primary artifact the user has to make educated operational decisions.
- **Capabilities**: Beyond the state, all resources export a certain degree of functionality and operative options. This functionality is decomposed via capabilities. A capability is a set of functionality that can be managed as a unit. Good examples of capabilities are routing protocol (OSPF and BGP are currently supported) for router resources or L2 circuits for a BoD domain resource. Capabilities also serve as management units to regulate user's access to resources.

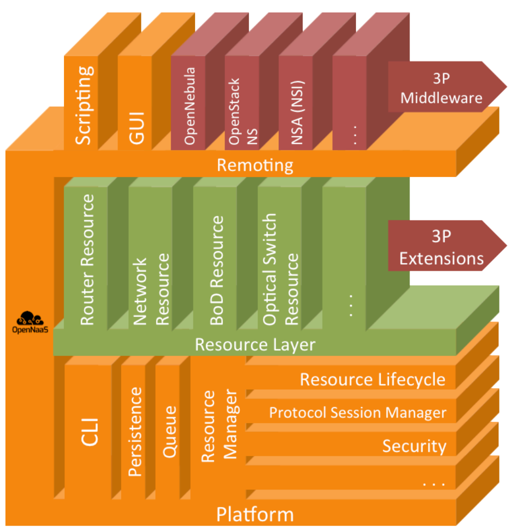

Figure 1 OpenNaaS framework layered architecture

All the resources have a defined lifecycle that regulate when they are reachable by users and the availability of the underlying model and capabilities. All resources available at a given OpenNaaS instance form the Resource Layer, in green on Figure 1. This is a network-specific and NaaS oriented hardware abstraction layer that the user can leverage to build his network intelligence applications and integrate with existing once. This network intelligence layer is pictured in red in Figure 1. The Resource Layer is shaped according to 3 key principles:

- User-triggered, fully automated and on demand operation over managed resources' capabilities.
- Recursive delegation of access and configuration rights over managed resources. A user must be able to delegate a subset of her access rights to another user, if allowed to.
- Lightweight abstracted operational model, decoupled from vendor specific details.

While the Resource Layer is simple enough to accommodate more sophisticated NaaS architecture and designs, they allow the creation of a reusable component toolset. These components provide several functionalities that are required in order to achieve an operational deployment. They can be decomposed as:

- Resource lifecycle management: handling resource initialization and cleanup, as well as subscription to resource states (i. e. a customer facing logical router).
- Security management: a key part of NaaS application is rights management. The framework includes an ACL engine to regulate this access to resources.
- Device configuration and integrity: operations on the Resource Layer must be translated to device-specific configuration in a way that respects the abstraction, other configurations and the institutional operation policies.
- User interface and Remoting: operations on the resource lifecycle and capabilities, as well as framework administration, are exported via a local (CLI) and remote interface (SOAP and REST).

Thus, using this framework as foundation, the Virtual CPE application will require a set of resources and capabilities. Table 1 below describes such capabilities, as well as their current implementation state. Capabilities marked with an asterisk are delegated to the customer.

IV. ENHANCING CPES

In this next section the Virtual CPE Scenarios are described. There are two scenarios, the Aggregated Virtual CPE and the Multiple Networks Virtual CPE. Although being two quite different setups, these scenarios show the similar Virtual CPE concepts and enabling technologies into play.

A. Aggregated Virtual CPE

A typical NREN in Europe has universities, secondary schools and research institutions as customers. Among those some of the universities are huge organizations while many of the other customers are small or medium seized. The role of an NREN is to be Infrastructure Provider as well as Service Provider. The largest of its customers may, on its turn, have the role of Service Provider as well as End User Community.

TABLE I
VIRTUAL CPE RESOURCES AND CAPABILITIES INVOLVED

Resource	Capability	Status
Router	Router virtualization	Implemented with logical routers.
	Interface virtualization	Implemented with 802.11q tagging.
	OSPF	Implemented.
	BGP	Not yet implemented.
	*Firewall	Not yet implemented.
IP Network	Topology and adjacencies.	Implemented using NDL.
	*OSPF	Implemented.
	BGP	Not implemented yet.
	*IPv4	Implemented
BoD	Dynamic L2 Circuits	Implemented via AutoBAHN for inter-domain GEANT circuits. BLUEnet will be supported for intra-domain circuits.

HEAnet is Ireland's NREN and are serving the full scale of institutions from large universities to primary schools. In Mantychore HEAnet is maintaining the Aggregated Virtual CPE scenario.

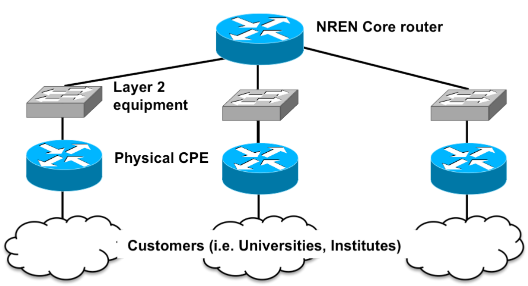

Figure 2 CPE in existing layout

The CPE is part of the customers network and typically interfacing the NRENs network by Border Gateway Protocol, BGP. Figure 2 show how the customers are connected to the core routers of the NREN. In this scenario (Aggregated Virtual CPE), the CPE-router is virtualized and localized at the NREN. Technology allows for operating several virtual routers within one piece of hardware. In figure 3 the virtual routers are illustrated as one Aggregated Virtual CPE.

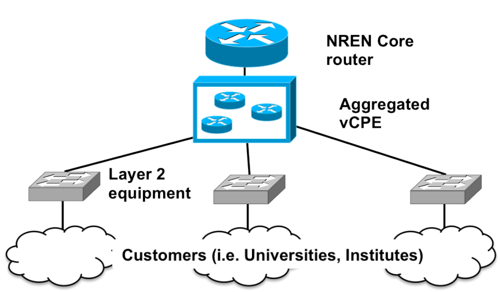

Figure 3 Aggregated Virtual CPE

The interface toward the client is flexible. Since the router becomes in fact a part of the client's network, it must be prepared to support a variety of features. The customer may use an Interior Gateway Protocol, IGP, on their existing network, so OpenNaaS must support a selection of IGPs. Open Shortest Path First (OSPF), Intermediate system to intermediate system (IS-IS) and Routing Information Protocol (RIP). Other requirements are Virtual Router Redundancy Protocol (VRRP) for facilitating redundant connections, Virtual Local Area Networks (VLANs) for facilitating a partitioned network. Network Address Translation (NAT) and Filtering on ports and addresses (firewalling). Those requirements are either implemented or on the road map to be.

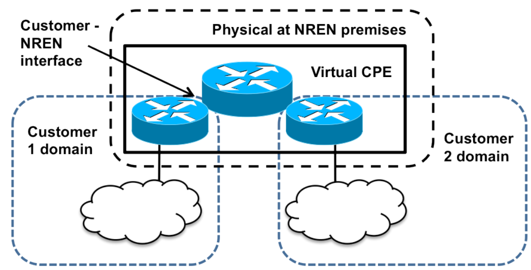

Figure 4 Network boundary

The boundary between the customer and the NREN is illustrated in figure 4. Customer domain is stretched up to the vCPE even it is localized central at the NRENs premises.

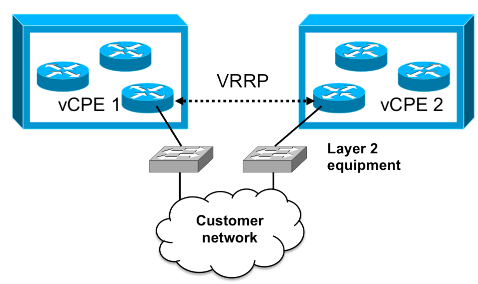

Figure 5 Redundancy ensures resilience

As illustrated in figure 5, duplication of lines towards a customer is possible for resilience. There are two main goals in vitalizing this part of the infrastructure. The first is to remove the necessity for a separate physical device (or pair of devices) to sit alongside an existing Layer 2 demarcation point, while still providing the flexibility of service that is needed toward the client.

The second goal is to delegate a level of control over the virtualized CPE to the user, one that they would not normally be able to achieve when their provider manages the device. This delegation is currently performed through the exposure of a REST interface. Integrity of the session toward the provider is not a primary goal here - providers already peer directly with clients' own routers in many cases, and have in place an infrastructure to ensure that misbehavior of any one client's router cannot affect other clients (e.g. BGP prefix filters.) However, the client must be able to activate a standard configuration for routing toward the provider - preferably without needing knowledge of the workings on BGP - and must be able to integrate the virtual CPE with their own network using RIP, OSPF and IS-IS. This implies a separation of functionality of the virtual device, so that different teams will be responsible for different parts of the configuration of the device. They must also be able to configure basic firewalling by means of access control lists. Using Mantychore OpenNaaS software for implementing Aggregated Virtual CPE will make it possible to add provisioning automation to the process, making the equipment configuration conform to NREN defined standards.

B. Multiple Networks CPE

The Danish NREN, Forskningsnettet, is connecting Danish universities and research institutions. A nation-wide Healthcare Data Network connects all health care and health research institutions. Today, these two networks use completely separate infrastructures as shown by the figure below.

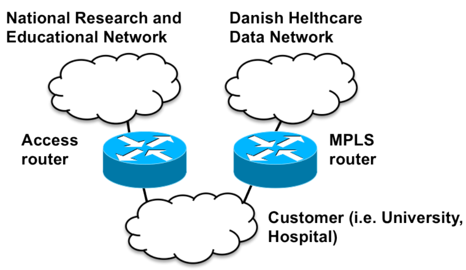

Figure 6 Separate attachments of two networks

UNI-C, a governmental institution, is providing connections for both networks. Several customers are connected to both networks utilizing different technical solutions and separate hardware. In the project use case UNI-C is using Mantychore software, OpenNaaS, for connecting customers to both networks trough same physical router. This utilizes virtual router instances instead of separate hardware. The router consolidation by means of virtual CPE is illustrated in the below figure.

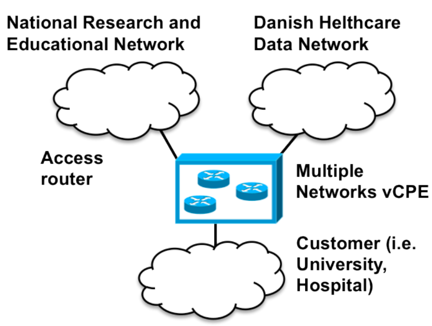

Figure 7 Multiple Networks Virtual CPE in context

The goal is to reduce the hardware required, by allowing two administrative teams to manage a single hardware device. At present, this demarcation can only be implemented by means of separate hardware, so a different router is needed for each network. With OpenNaaS, it is intended to use the administrative partitioning that can be provided by logical networks to provide access to both networks, managed separately, from the one router.

Deployment of the Multiple Networks vCPE will possibly be done related to new customers and for existing customers in the peace of requirements for upgrade of capacity or equipment. Provisioning automation and conformant configuration are asserts like in the Aggregated Virtual CPE scenario. In this scenario the governmental institution is Infrastructure provider while the two network organizations are Service Providers.

C. Qualitative analysis

The main benefit of Virtual CPE scenarios is the reduction on cost of extraneous equipment at the border between the NREN and the client. This is possible where this equipment is needed to provide a logical demarcation but not a physical conversion. As shown, this is accomplished by the virtualization-enabled possibility of having several customers to consume virtualized instances on the same physical equipment. The fact that this physical equipment is not only concentrated by also co-located allows for several secondary but relevant benefits on operational expenses. This allows for reducing space, power and cooling needed due to physical equipment, and therefore the cost and environmental impact of the same.

Although the initial setup requires an investment that cannot be justified to serve only one user, it quickly scales to considerable saving once more users are plugged the Virtual CPE. This setup also benefits for gradual deployment to existing or new clients, easing both the pilot phase and the production operations deployment. The Mantychore project has made a projected savings evaluation on the planned setups. Here you can find the detail of the savings that are expected on the pilot phase. The chosen setup is based on Juniper MX-240 routers. This models support up to 16 logical routers. In this setup, two of them are reserved for control and resilience connectivity for this and a backup router. We can then server router pairs (one for the customer and one for the NREN) to 7 customers per router. The setup is duplicated in a secondary physical router for backup purposes. This setup is then compared to traditional CPE based both on Cisco and Juniper equipment:

TABLE 2
CPE SETUP COMPARISON

	Virtual CPE MX 240	Multiple CPE ASR1002-F	Multiple CPE MX 10
Manufacturer	Juniper	Cisco	Juniper
# Users Supported	7	7	7
NREN CPE	Yes	Yes	Yes
Space	5 U	14U	14U
Power (kw)	2000	3290	2632
Weight (kg)	58.97	127	96

Below there is a projection of the savings for equivalent setups that fit 7 customers:

TABLE 3
VIRTUAL CPE SETUP SAVINGS FOR 7 CUSTOMERS

	MX 240	ASR1002-F	MX 10
Manufacturer	Juniper	Cisco	Juniper
Space Saving	-	560%	560%
Power Saving	-	329%	263%

Weight Saving	-	430%	325%

Beyond this expected savings in space, power and weight there is also a marginal benefit on number of interfaces needed, which are also a source of cost and often a scalability-limiting factor. Although equipment cost can vary largely due to external factors, the cost analysis performed at the NRENs projects a minimum saving of 30% and an expected 50% in equipment purchase. Additionally, greater flexibility in network services that can be provided by the NREN to its clients, including many-point connections and short-term elastic services, are enabled by a virtual CPE setup. This can be, at later stages, be the foundation of other NaaS services. Provisioning automation is an additional gain as it minimizes the source of diffuse errors and possibly boosts productivity in administration.

However, virtualization concepts are not supported with equal degree in common monitoring and network automation software. This can introduce incompatibilities that incur on hidden costs and complexities. A good example of this is humans and or applications that require direct CLI access to devices. While CLI provides a logical router abstraction, this is not complete and its feasibility must be tested under real world scenarios. This and other complexities will be further explored during the pilot phase.

V. Future Work

In the mid-term, these two scenarios will be deployed for trial in the mentioned NRENs, HEAnet and UNI-C. The first will focus on the Aggregated Virtual CPE scenario while the second will focus on the Multiple Networks CPE one. These deployments will be coordinated from the Mantychore FP7 project, with an eye on assessing their viability for production environment. All deployments will include real end users.

The Aggregated Virtual CPE will be built on top of Juniper M-series or MX-series routers, connecting to Cisco (CRS-1 or CRS-8) routers for onwards connectivity to the Internet. The internal domain connectivity will be provided by a Point-to-point Ethernet service, providing Layer 2 port and VLAN transmission (actual implementation is arbitrary, but will be tested on Cisco 7600 MPLS/ME3400 switched Ethernet network).

The Multiple Networks Virtual CPE will be built on Juniper MX-series routers interfacing MPLS to the one side and wavelength connections to the other. A successful pilot will open the door to further, more sophisticated NaaS offerings and capabilities from the involved NRENs. Some of these are already being researched inside the Mantychore project although their description goes beyond the scope of this paper. Main research efforts in the Mantychore FP7 project are energy infrastructure consumption awareness and infrastructure marketplace.

VI. Conclusion

This paper has described how state of the art concepts regarding virtualization can be applied to a concrete use case, Virtual CPE, in order to leverage cost and operative benefits and improvements. A tool that implements this concepts has been described, OpenNaaS. In summary, it allows to get resources from the network infrastructure: routers, switches, links or provisioning systems; abstract them to service resources, independently of vendor and model details; embed instantiation of virtualized resources in the regular BSS/OSS workflows, and delegate management permissions over the infrastructure resources they own so that "Infrastructure integrators" can control them during a period of time. The application of these capabilities to the actual CPE deployments allow for the enhanced scenarios described. The main benefits described can be categorized into capital and operational expenditure. Not only a Virtual CPE setup allows to server customers with a less expensive equipment setup, but several benefits in space, power and operations steam also a significant saving.

Acknowledgment

The work in this paper has been funded by the European Union through the FP7 project MANTYCHORE (INFRA 26157).

References

[1] Mantychore website. Available online at: http://wwww.mantychore.eu/

[2] 4WARD: Introduction to a 4WARD Architectural Approach (D-0.4), URL: http://www.4ward-project.eu/index.php?s=Deliverables

[3] 4WARD website. Available online at: http://www.4ward-project.eu/

[4] GEYSERS website. Available online at http://www.geysers.eu/

[5] M. Campanella, "Federated E-Infrastructure Dedicated to European Researchers Innovating in Computing Network Architectures," Future Internet Conf., Apr. 2, 2008; http://www.fp7-federica.eu/

[6] GEANT website. Available online athttp://www.geant.net/pages/home.aspx

[7] OpenNaaS website, Available online at: http://www.opennaas.org/

[8] L. Peterson, A. Bavier, M. E. Fiuczynski, and S. Muir., "Experiences Building PlanetLab," Princeton University, Technical Report, Nov 2006.

[9] New Generation Network Architecture, AKARI Conceptual Design (Ver.2.0), NICT, May 2010 (available at http://akari-project.nict.go.jp)

[10] I. Houidi, W. Louati, D. Zeghlache and S. Baucke "Virtual Resource Description and Clustering for Virtual Network Discovery". Proc. IEEE International Conference on Communications Workshops ICC 2009.

[11] Jorge Carapinha and Javier Jim "Network virtualization: a view from the bottom". In Proceedings of the 1st ACM workshop on Virtualized infrastructure systems and architectures (VISA '09).

[12] G. Schaffrath, C. Werle, P. Papadimitriou, A. Feldmann, R. Bless, A. Greenhalgh, A. Wundsam, M. Kind, O. Maennel and L. Mathy."Network virtualization architecture: proposal and initial prototype" In Proceedings of the 1st ACM workshop on Virtualized infrastructure systems and architectures (VISA '09)

[13] I. Houidi, W. Louati, W. Ben Ameur, D. Zeghlache, "Virtual network provisioning across multiple substrate networks", Computer Networks, Volume 55, Issue 4, 10 March 2011, Pages 1011-1023