HEAD IN THE CLOUD, FEET ON THE GROUND: VIRTUALIZATION OF THE RESIDENTIAL GATEWAY

Authors: Barak Hermesh, Shaul Shulman, Guy Ray, Paul Mannion, Eoin Walsh
Intel Corporation

1 Abstract

The benefits of virtualizing the functionality of the broadband home gateway into the cloud have been discussed for some time in the telecommunication industry. These include OpEx reduction through ease of home gateway SW deployment, faster time-to-market of new features, and ease of service-chaining with little dependency on capabilities of the residential HW resources.

As we enter the DOCSIS 3.1 era, the transition to an end-to-end virtual gateway architecture not only is needed but also technologically attractive.

This paper describes the Virtual Home Gateway using Software Defined Network principles. We will discuss a reference architecture and present analysis of benefits in terms of easier software maintenance and reduction in time to market of new services. We will also discuss some of the implementation challenges in the path of making the Virtual Home Gateway mature for mass deployment: open networking while preserving subscriber privacy and intelligent hair-pinning over the wide-area network under constrained uplink resources.

2 INTRODUCTION

Transformation of the networking architecture into a Software Defined Network (SDN) seems no longer to be a question of “if,” but a question of “how” and “when.” Successful implementation of SDN in the Data Center space has triggered interest to apply the same principals to the telecommunications network due to the many similarities in both the challenges that SDN can overcome and in the methods that can be reused or replicated in telecommunications and broadband residential access networks. The ecosystem of SDN based solutions is growing rapidly and is available for MSOs. These solutions include Commercial Off-The-Shelf (COTS) hardware platforms and software tools. Since these are in general applicable for any networking application, a strong business incentive drives a very high speed of innovation that any particular industry can benefit from.

In the broadband residential access industry, two major growth vectors have been present since the early days of the introduction of Internet connectivity to the home. These are the connection speed race and the growth in the diversity of the services provided over the data connection. Both have expanded rapidly over the years, presenting a challenge of scalability of the infrastructure to continue supporting this growth. In the case of the speed of access, a new DOCSIS standard (DOCSIS 3.1) was recently introduced. It allows a much more efficient use of spectrum, enabling scalability of the residential connection data rates into the Gigabits per second (Gbps) range over the existing Hybrid Fiber Coax (HFC) infrastructure. In parallel, introduction of Fiber to the Home (FTTH) technologies easily allows multi Gbps speeds. These two advances create a future proof, scalable physical layer infrastructure, ready to accommodate anticipated as well as unanticipated demands in communication bandwidth.

On the side of the services however, nothing has significantly changed in terms of infrastructure. Conceptually, the network architecture for delivery of Gateway services looks the same as it looked 15 years ago. Functionality of the Gateway is embedded into the GW's hardware, firmware and software on a multitude of proprietary platforms from different vendors. Adding a new service requires a field upgrade of software and firmware on the embedded CPE device (different update per platform type). This severely limits the TTM of deployment and raises cost, putting a burden on operations and risking future scalability. An anecdotal indication of complexity growth of the embedded gateway is the fact that its SW stack image has grown about 100 times in size since the early days of a Cable Modem.

As a parallel process, over the top (OTT) providers are pushing forward to enable new services which are not dependent on the home user platform. The MSO's control of the whole MSO network and the CPE is not utilized to its full potential under existing networking architecture.

SDN and Network Function Virtualization (NFV) technologies present a wonderful opportunity for MSOs to upgrade the L2-L3 networking infrastructure into a more agile, scalable and future proof architecture. Virtualization of the network functions is also symbiotic with existing trends, including "fiber deeper," digitization of the last mile access from analog fiber to digital fiber, a move to distributed architectures of the Cable Modem Termination System (CMTS), and transformation of linear video into IP video. With virtualization, Service Providers and Operators also can offer services never before possible or not practical to implement, such as those based on "hair-pinning" and those that require intensive data path computation.

Application of SDN and NFV into Cable MSO networks has been talked about in the past several years ([1],[2]). As the ecosystem matures it seems we are near the time of moving from discussion into implementation stage. The purpose of this paper is to expand on the architecture, review and compare different architecture alternatives and offer our view of what the virtual CPE will look like and how commercial service providers and the surrounding ecosystem can benefit from the new architecture.

The paper is structured as follows. Part 3 sets a common terminology. Part 4-7 describes a proposed virtualized architecture for the Gateway and discusses the advantages of the architecture. Part 8 further articulates the advantages of the virtual CPE, Part 9 discusses the challenges and a summary follows in Part 10.

3 TERMS AND DEFINITIONS

As the industry has not converged on terminology of the network elements of the future architecture and in order to avoid ambiguity, the following terms are proposed and will be used as defined here in the rest of this paper.

pCPE – A physical CPE. The term describes the actual hardware residing in the subscriber's premise, which could be residential or business subscriber.

vCPE – A virtual CPE. This is a logical entity providing services that were traditionally provided by the pCPE. A vCPE is a software entity residing in the provider edge (PE), Place of Presence (PoP), data center or a mixture of the locations.

Home Gateway – The collection of functions that provides WAN (broadband) connectivity, routing, and bridging into the home and across the wireless and wireline networks in the customer premise.

Embedded Home Gateway – A home gateway whose execution is done exclusively on a pCPE

Virtual Home Gateway – A distributed network solution implementing the home gateway functions across a pCPE and a vCPE, in which the pCPE is a slave to networking decisions made by service logic executed on the on the vCPE.

vCPE Home – the vCPE component of a **Virtual Home Gateway** solution

PoP –A regional Place of Presence on the network infrastructure.

Customer Edge (CE) – The topological location of the customer facing edge of the MSO network. For a cable network, this would typically be where the pCPE is.

Provider Edge (PE) - The topological location of the provider's network edge, facing the access network. For a cable network, this would typically be the edge of the MSO MPLS network facing one or more CMTS/CCAPs. Some or all vCPE functions are hosted by a server hub located at the PE.

4 NETWORK ARCHITECTURE

In this section we analyze and compare several possible network architectures. We compare the architectures using the following parameters:

Efficiency – how efficient is the architecture in terms of cost of equipment.

Agility – how complicated it is to introduce new services or to upgrade services.

Operational expenses – how are the operational expenses affected by an architecture compared to alternative architectures.

4.1 MSO Network Topology

A cable MSO network consists of many elements. This section sets the terms of the elements relevant for the rest of the paper, illustrated in Figure 4-1.

The CE or Customer Edge is the edge of the MSO network residing in the customer premises. The CE equipment consists of a Cable Modem with potential Embedded Service Functional Entity (ESAFE) such as an embedded router.

The CMTS or CCAP terminates the HFC network and provides access to the MSO network.

The edge of the provider core network which is usually MPLS/Metro Ethernet is the location where the access network meets the MSO core network. The Provider Edge (PE) is usually located within the vicinity of the access termination equipment and usually covers a few CMTS/CCAP nodes. Data path processing related functions of the virtual home GW usually reside at the PE hub. Other functions may reside further away. The PE is sometimes known in the industry as the Central Office (CO) but in this paper we use the term PE.

The Point Of Presence (PoP) is a regional center aggregating traffic from multiple CE entities. Typically, a country or state has one or a few PoP centers. A PoP usually employs cloud technologies.

A service provider data center is usually a single location in one geographic region providing functions for the whole service provider network. The data center usually employs large-scale cloud technologies.

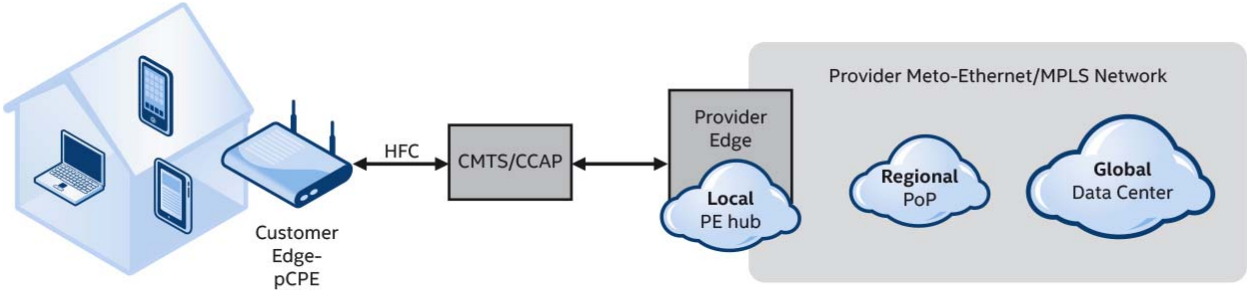

Figure 4-1 MSO Network Topology

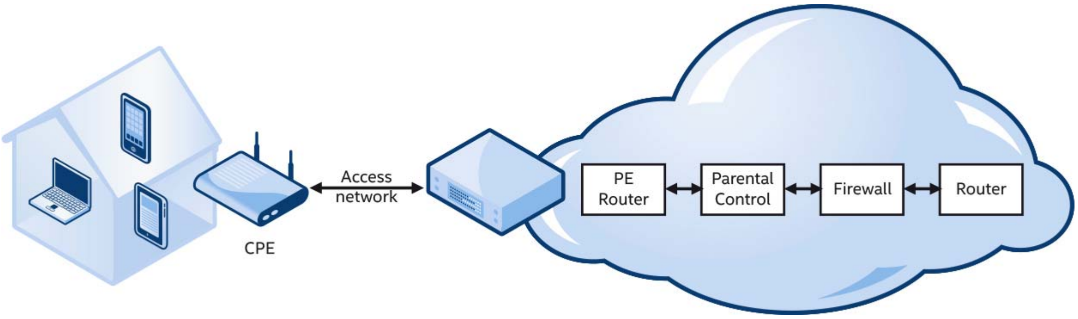

Figure 4-2 Legacy Service Provider Network

4.2 Existing Architecture

In today's existing architecture, the Home Gateway is an Embedded Home Gateway executing all gateway functionality locally. Traffic between devices in the home and the Internet flows through and is processed in the Home Gateway. It is then forwarded through the access network to the remote termination point (CMTS, DSLAM, OLT or other) and goes to the service provider's network. The service provider deploys carrier grade services like firewall and parental control at the provider's network. In a non-virtualized network, these services are implemented in proprietary blue-box equipment. Traffic flows through the network equipment where it is being processed.

Figure 4-2 illustrates the existing non-virtualized service provider network with parental control and firewall functions as examples. A typical home with connected devices contains a single embedded Home Gateway connected to the access network termination point. The service provider network contains blue-box network equipment implementing services (also known as network functions).

The existing architecture is efficient in the short term as proprietary equipment implementing a single network function is usually cost optimized. However, upgrade, enhancement and replacement of network functions has high cost overhead.

Introduction of new services becomes very complicated with the existing architecture, mainly due to the large variety of proprietary network equipment. A cable operator has on average about six different pCPE hardware platforms deployed, coming from different OEMs. Adding an embedded network function requires at least a software upgrade to all pCPE devices, assuming the pCPE hardware can even support the new function.

One lesson from the past illustrating the complexity is DOCSIS multicast forwarding:

DOCSIS 3.0 introduced a technique called Multicast DSID Forwarding (MDF) [3]. Its purpose is to allow for the CMTS to dictate the multicast forwarding at the CM. This was done, in part, because of the huge complexity involved with upgrading all CM (pCPE) devices to support IGMPv3. In a sense, MDF follows SDN concepts, where the forwarding device (CM) is a slave and the actual logic is determined by the software residing on a remote device (the CMTS).

A report by Analysis Mason [4] shows that moving customer households from an embedded Home Gateway to a fully virtual one reduces operational expenses by 64%. The two most significant contributors to operational expenses are customer visits and service calls. Both are much higher with the existing architecture, compared to a fully virtualized environment.

4.3 The Fully Virtualized Network

The Home Gateway in a virtualized environment is split between the pCPE and the vCPE. The pCPE owns the physical LAN ports and the actual packet forwarding in the home domain, both in-home and through the service provider.

In a virtualized environment, the physical elements are reduced to provide physical layer (Phy), Media Access Control (MAC) connectivity and data forwarding. The actual control and decision making is done by external software running on a virtual platform. The pCPE in this case would serve as a SDN switch exposing LAN and WAN ports to the SDN controller. The Home Gateway logic is virtualized by the vCPE-Home software running as a collection of virtual network functions in the service provider network.

The vCPE performs traditional Home Gateway functions like configuration of the LAN IP network by DHCP services, NAT, IP routing and more. For each pCPE, an instance of a vCPE is generated by the orchestration upon pCPE initialization. Each CPE obtains an individual configuration and state maintained mostly by the vCPE.

The vCPE is the management endpoint for the Home Gateway. As such, it exposes management interfaces like TR069 and web UI towards the home and the service provider. As the owner of the home GW management and configuration, the vCPE manages the configuration of the pCPE. This includes querying for counters and statistics but also controlling the port configuration (like WiFi SSID configuration).

A network orchestrator is software running on a virtual platform. It controls the instantiations and monitors the lifecycle of the Virtual Network Functions (VNFs) at the operator network including the vCPE. The orchestrator also communicates with the SDN controller, feeding it with forwarding policies based on the VNFs and on the service chaining policies.

In a fully virtualized network, the network access equipment (CMTS, DSLAM, OLT) is also virtualized. A CMTS, for instance, would have its forwarding logic reduced to a SDN switch and the control would come from the SDN controller. It is however expected that at first stage the access network would remain as today and the traffic from the pCPEs to the operator network would be tunneled over the underlay access network as explained in 4-6.

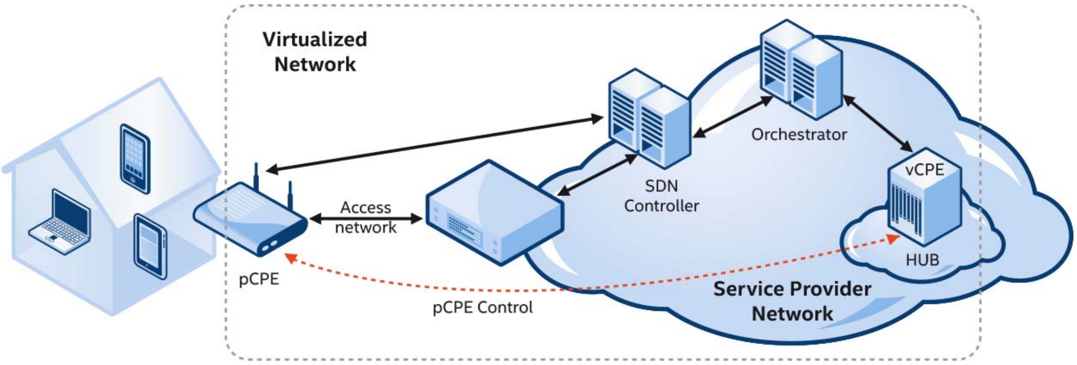

Figure 4-3 The Virtualized Network

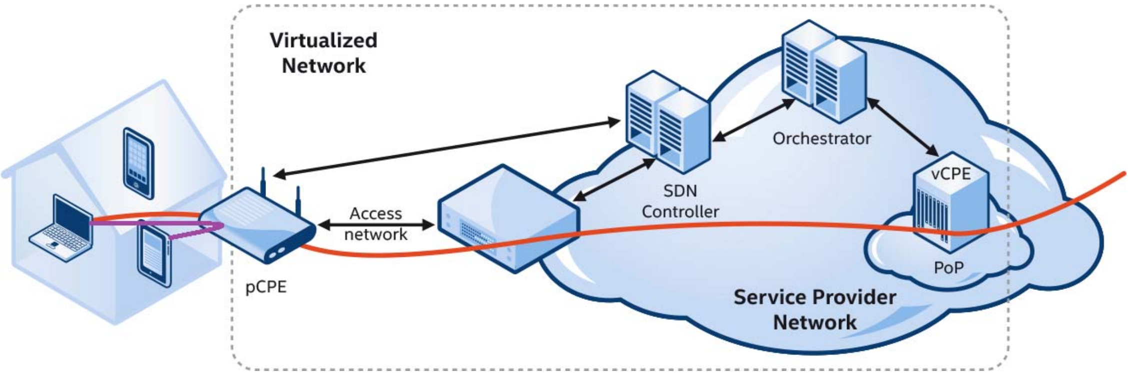

Figure 4-4 Straightforward pCPE Forwarding

Figure 4-3 illustrates the major elements of the virtualized network.

4.4 Software Defined Forwarding at the pCPE

The physical CPE performs packet forwarding. The rules by which the pCPE forwards the packets are set by the SDN controller. It is expected that in conventional cases traffic on the same LAN segment will be simply forwarded from the source LAN port to the destination and that traffic to an external destination will be forwarded to the access network termination point. Due to the power of SDN, exceptions could be made.

Figure 4-4 shows two forwarding paths. One LAN-WAN marked in orange and one LAN-LAN marked in violet.

The flow would start with a pCPE having pre-configured forwarding rules set by the SDN controller (this is called proactive mode). The rules would say roughly:

1. For any unknown packet, consult with the SDN controller.

2. For all packets whose destination is not in the home network, forward the packet to the WAN port (potentially with tunneling)

3. Specific rules based on destination MAC address for L2 forwarding in the home, set by dynamic consultation with the SDN controller

Since the forwarding logic at the pCPE is dictated by the SDN controller, topology changes and exceptions could be easily made and deployed. One novel example is hair-pinning.

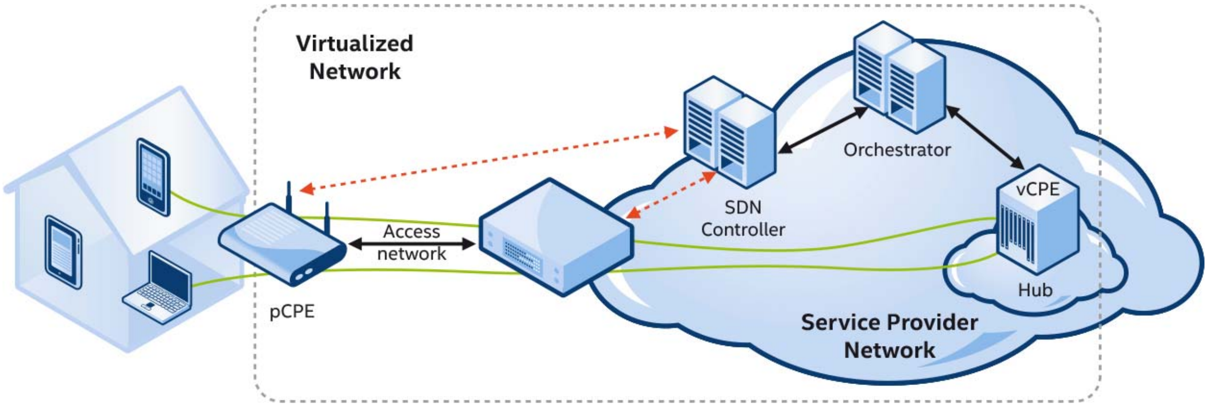

Figure 4-5 Hair-Pinning

Hair-pinning, in this context, is the concept of sending a stream of packets outside of the home network for processing at the vCPE and then looping it back in to the target port. Figure 4-5 shows a case of hair-pinning between a laptop computer and a smart phone. Although both are on the same network, an administrative decision had been made to hair-pin the traffic through the vCPE. Thanks to the flexibility of SDN, hair-pinning could be done for selected flows or even selected packets.

There are multiple usages for hair-pinning, starting with enhanced in-home network security and going through remote diagnostics and troubleshooting. At this point in time, since the broadband access network is asymmetrical, hair-pinning is limited by upstream capacity. Thus, it should be applied with care.

4.4.1 Network Security via Hair-Pinning

In the modern home, the assumption that in-home traffic is safe is no longer valid. Malware can be injected into the home via a guest joining the home network, an attacker connecting to the home wireless network and in other ways. As the home network becomes more essential, it becomes mandatory that traffic originated in the home and terminated in the home is subject to monitoring by Intrusion Prevention Systems (IPS).

Running a full blown IPS on the pCPE is not feasible today and will not become feasible in the future as the computation power available for applications grows at a rate not higher than home network throughputs. In other words, the gap between the computation the pCPE can perform at a reasonable cost and the computation required by IPS does not close. The only feasible approach is therefore to take advantage of the elasticity and the power of the cloud and run IPS as a VNF. The SDN policy becomes such that when required (and no more), selective sessions get hair-pinned through the IPS for inspection.

Imagine a case where the alarm system starts communicating with the smart TV. This might be considered an irregular pattern and will then be hair-pinned for inspection. After a few packets, the IPS, being a logical part of the vCPE, may conclude there is no threat and instruct the SDN controller to terminate the hair-pin. The SDN controller will then reconfigure the pCPE SDN forwarding logic to switch the traffic locally.

4.4.2 Real Time Troubleshooting via Hair-Pinning

From time to time, subscribers experience in-home network connectivity issues. Since MSOs are responsible for home networking, such service calls frequently end up with technician visits generating high expenses to the MSO.

With SDN it is possible to troubleshoot such issues remotely. The technical support team sets up hair-pinning to all of the subscriber's in-home traffic, or to selective flows if identified as faulty. Using standard monitoring and analysis tools, the technician may detect the issue and apply corrective measures.

Other advantages of supporting SDN at the pCPE include flexible residential topology setup, simple creation and teardown of tunnels and future proofing.

4.5 The Virtualized Hub

An alternative to a fully virtualized network is a hybrid where the existing gateway services are embedded in the pCPE and additional services are located in a hub at the MSO core network. The MSO core network is SDN operated and network functions are VNFs.

A hub is given here as a generic term for the cloud compute cluster where VNFs run. A virtualized hub operating with a pCPE that is SDN unaware hosts VNFs at the hub without extending the software defined network into the customer premises. The pCPE in this case contains an Embedded Home Gateway with multiple WAN-facing tunneling interfaces. Classifiers provisioned to the pCPE direct traffic into tunnels based on the VNF service chain required to operate on the traffic. Figure 4-6 shows an example of three tunnels representing three classes of traffic. The pCPE performs some GW functions locally while other services are executed at the hub.

Compared to the existing architecture, this architecture's operational expenses for the VNFs are reduced, while the same level of complexity and operational expenses at the pCPE is maintained.

Service agility is achieved by deploying and upgrading VNFs. Orchestration includes setting up tunnels, tunnel terminators are service chains. Since the pCPE is not a part of the virtualized network, agility of traffic forwarding is limited to the service provider network. Flows like hair-pinning are not possible as LAN to LAN traffic is always terminated at the pCPE. pCPE related OpEx and complexity remain the same as in the existing architecture.

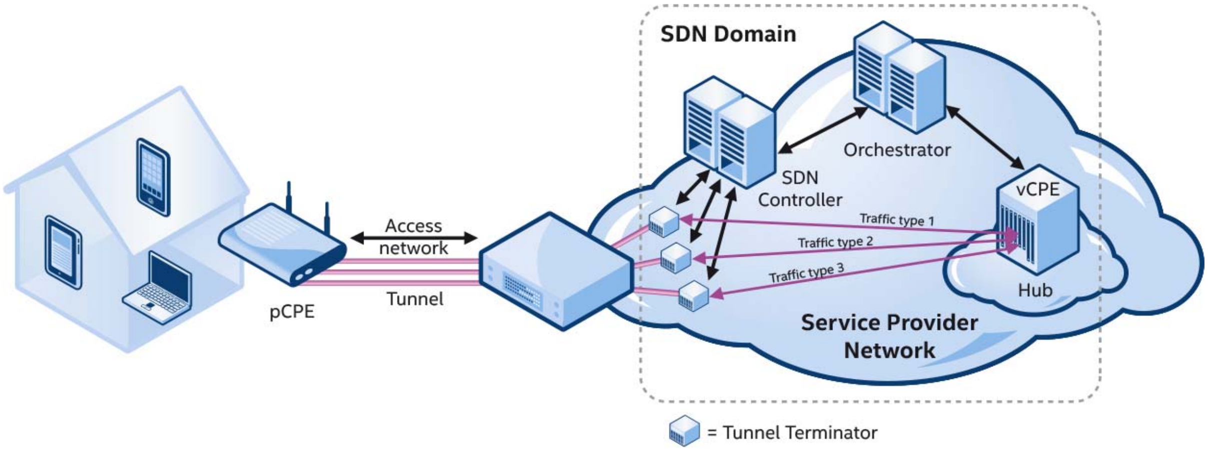

Figure 4-6 The Virtualized PoP

4.6 The Hybrid Network: Overlaying Non-Compliant Network Equipment

It is expected that not all network equipment will become SDN compliant at once, there will be gradual transition from the existing architecture to a fully SDN compliant architecture. In order to traverse through non-SDN network equipment, tunnels are created as bridges between SDN domains. Tunnels are logical point to point connections on top of an underlying network. Popular tunneling techniques include Layer 2 General Routing Encapsulation (L2GRE also called EoGRE) over IP and Virtual Extensible LAN (VXLAN) over UDP.

The most probable scenario is that while the home and the service provider core network become SDN compliant, the access network is still not SDN compliant. In this case, a tunnel is used to connect the residential SDN domain and the service provider SDN domain. The pCPE is located at one end of the tunnel and a dedicated tunnel termination network function is located in the service provider network and serves as other end of the tunnel.

It is allowed that multiple tunnels exist for the purpose of load balancing, redundancy and service segregation. The SDN controller instructs the pCPE and the tunnel terminator to direct traffic into a selected tunnel using standard SDN rules.

Figure 4-7 shows a case of a single tunnel from one pCPE to a tunnel terminator. Note that the SDN controller controls both the pCPE and the tunnel terminator. The tunnel bridges between to SDN islands comprising a single SDN domain together.

4.7 Tunneling Overhead

One of the drawbacks of tunneling is that it adds network overhead. L2GRE [5]over IPv4 overhead is 14+20+16 bytes for plain Ethernet header+IPv4+GRE. 50 bytes appended before each L2 frame, assuming an IMIX [6] average packet of 700 bytes, GRE tunneling adds an overhead of 7% resulting in a drop of 7% of effective data throughput.

5 THE pCPE ARCHITECTURE

The pCPE consists of four layers. The physical ports, the SDN forwarder, the forwarder control and the provisioning/OSS layer.

The bottom layer is that of the physical ports. All network ports are exposed to the layer above.

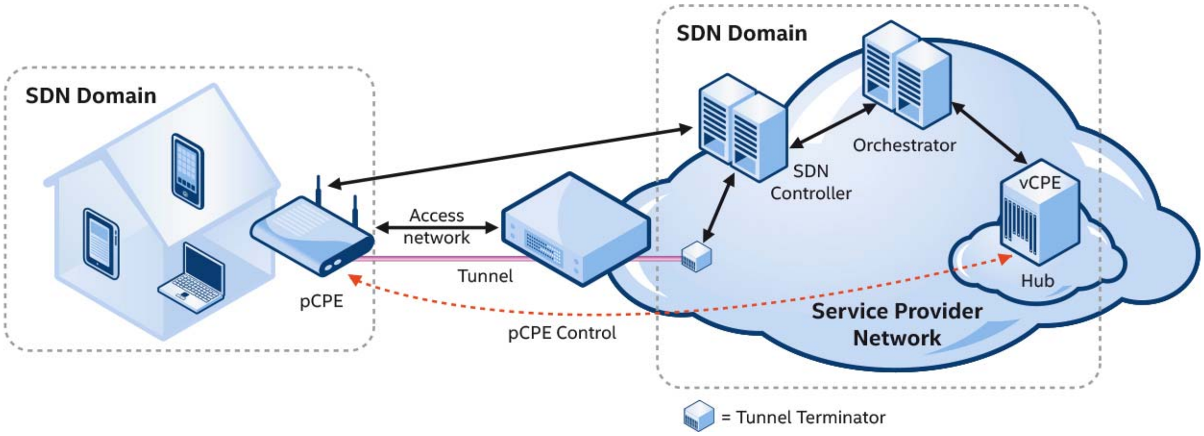

Figure 4-7 Transition Phase via Overlay

The layer above the ports is the SDN rule based forwarding layer. The forwarding layer receives packets from the ports and forwards them to their destination ports based on the SDN rules dictated by the SDN controller. One possible implementation of this layer is the Open VSwitch forwarder [7].

The layer above the SDN forwarder is the forwarding control layer which consists of two mutually exclusive logics. The SDN protocol termination point comes to action when the link to the vCPE is established. In this case, all forwarding rules are set by the SDN controller by sending them to the protocol termination point. In the case of Open Flow, this would be an Open Flow protocol stack endpoint.

The offline mode logic comes to action when the link with the vCPE is down. This logic makes sure that LAN to LAN forwarding is made possible by setting L2 forwarding rules to the forwarder and by providing L3 configuration to the hosts on the LAN if required.

The top layer consists of the device management, provisioning and state monitoring.

Provisioning starts when WAN connectivity is established and includes notifying the orchestration logic that a pCPE is up, pairing with the vCPE and obtaining runtime configuration.

The device management module is the local endpoint of the vCPE-pCPE device management. The vCPE configures the pCPE runtime parameters by utilizing this module. The vCPE also queries the pCPE device status and statistics by communicating with this module.

The link state monitor is a daemon monitoring the connectivity between the pCPE and the vCPE periodically. If the link goes down, the monitor alerts the provisioning module which may change the operational state from SDN forwarding to offline forwarding until connectivity is re-established.

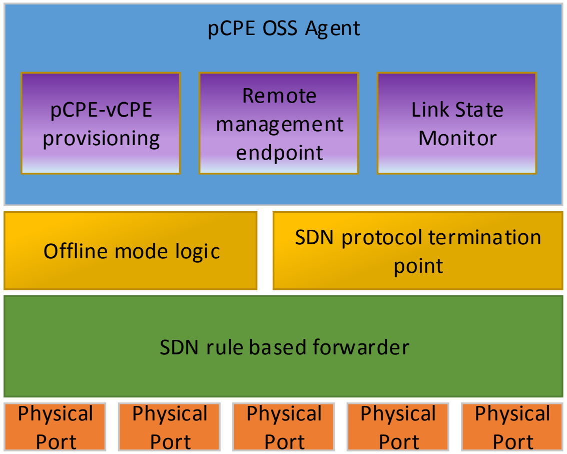

Figure 5-1 pCPE Architecture

6 THE vCPE ARCHITECTURE

The vCPE is a set of software services executed on virtualized platforms located at the PE and sometimes also in other locations. Unlike the pCPE, the vCPE can be implemented in various ways starting with each vCPE being a completely independent self-contained micro service, containing all vCPE functionality for a single home GW and going to a highly distributed architecture where single tenant network functions support multiple pCPEs and service chaining is applied in order to link all network functions together.

The vCPE provides the following services –

- Routing with address translation (NAT)
- Port forwarding, UPnP, port triggering
- Firewall
- IP configuration for hosts on the LAN
- DNS
- Advanced services like parental control, centralized storage (NAS), UPnP media server, DLNA media server and more

In this section we discuss and analyze three possible vCPE architectures. The key difference between the architectures is in the grouping of functionalities into network functions. The tradeoffs are simplicity vs. performance and scale.

First, we provide an observation as to the relation between the MSO oversubscription ratio and the computational power of the vCPE server.

All architectures analyzed here comply with ETSI ***Network Functions Virtualization (NFV); Architectural Framework*** [8]

6.1 vCPE Data Path Computation

All MSOs over-subscribe their subscribers. If the physical network infrastructure is capable of carrying a given bandwidth, the sum of the bandwidth provisioned to all subscribers would be much higher. The ratio between overall subscribed bandwidth and available bandwidth is defined as the oversubscription ratio. The oversubscription ratio is carefully set by the MSOs to meet the statistics of bandwidth consumption of the subscribers, relying on the spread of bandwidth demand (across time) among a large number of subscribers.

Since each unit of data going through the access network is also processed by a vCPE, the overall processing power required for all vCPEs is linear with the physical network infrastructure bandwidth. The compute load of the vCPE server platform benefits from oversubscription.

As an example, a node of 100 subscribers, each provisioned with 100Mbit/Sec is connected to the MSO network over a 1Gbit/Sec link. The oversubscription ratio here is then [100 x 100Mbit] / 1Gbit = 10. A server capable of processing 10Gbit/Sec of CPE network traffic at the Provider Edge can support 10,000 subscribers. Each is provisioned with 100Mbit/Sec given the oversubscription ratio of 10.

6.2 The vCPE as a Micro-Service

Figure Figure 6-1 shows multiple instances of vCPEs, each instantiated as a micro service on a virtual machine. The number of services is dictated by the number of pCPEs currently provisioned, but the number of virtual machines and the allocation of a vCPE service to a virtual machine is dictated by the orchestration layers and is based on optimization and load balancing. The micro services are expected to run within a virtual isolated environment such as a ***Linux Container*** [9]. Note that a real world vCPE is expected to have more services than illustrated in the figures below. The figures below show a few services just for illustration.

Each micro service is a fully self-contained vCPE. Each micro service is very similar to an embedded home GW module. This architecture is simple to orchestrate and offers the simplest development model. However, from the performance aspect this solution is the least efficient. Efficiency is low since packet forwarding, which is the most intensive computational task of the vCPE, is done in many different contexts resulting in context switching and cache thrashing.

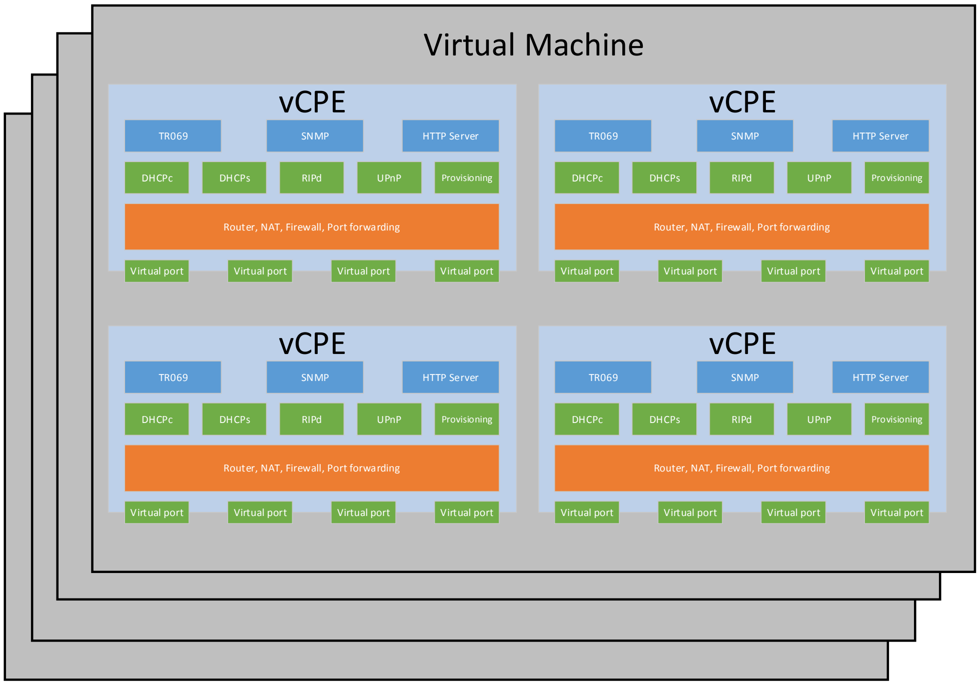

Figure 6-1 vCPEs as micro services

6.3 Split Control and Data Path

Another alternative that is expected to be more performance-efficient is to separate the data processing from the vCPE control. The vCPE control is instantiated per pCPE as in the previous case. However, the data path processing including routing, address translation, port forwarding, firewall and other network functions is done on optimized dedicated single tenant virtual network functions. The advantage of this approach is that it is more efficient both in resource consumption and in compute resources to run all data through a single VNF than to run it through many smaller micro services.

The VNFs are CPE aware so there is a logical context per CPE. The vCPE control installs configuration rules for the centralized VNFs (router, firewall etc).

As an example, if the vCPE-Control TR069 agent is instructed to set up a new firewall rule, it configures the firewall VNF to add the rule for the CPE it represents. The SDN rules are set such that control traffic from the pCPE is forwarded to the vCPE-Control entity where other traffic is forwarded to the service chain containing the NAT, router, firewall and other VNFs.

Figure 6-2 shows multiple vCPE-Control micro services and two centralized VNFs – one for routing and address translation and one for firewall.

This approach is more performance optimal but it comes with the extra complexity induced by the need to orchestrate multiple functions per pCPE and by having a control and communication channel between the vCPE-control and the data path VNFs.

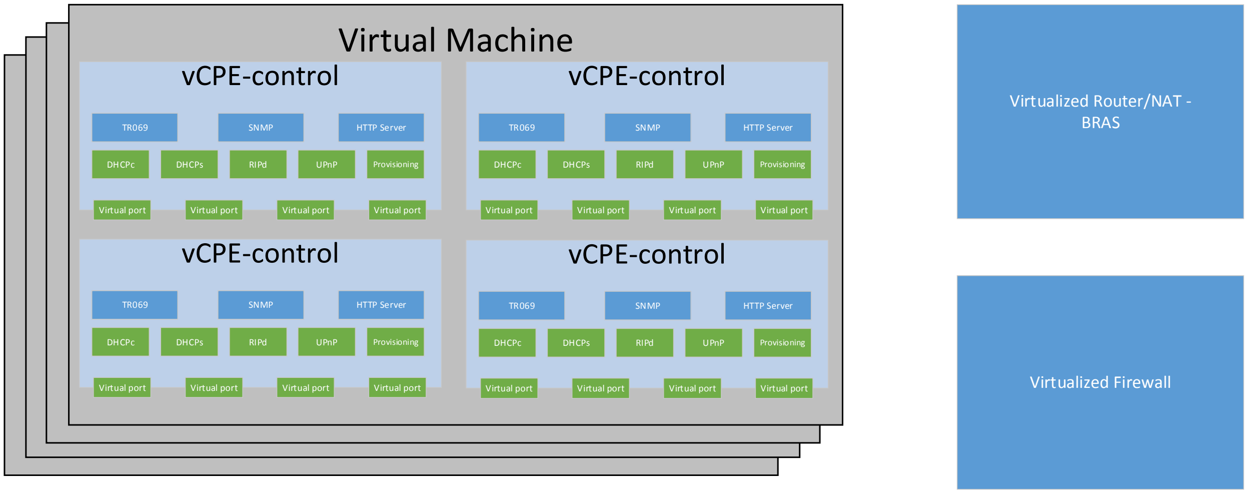

Figure 6-2 vCPEs with data path functions are centralized entities

6.4 Centralized Control and Data Path

An even more centralized alternative runs some of the control entities as centralized VNFs. Two candidates are the vCPE DHCP server and DHCP client. Instead of instantiating one per pCPE, two centralized, multi context VNFs are instantiated. The DHCP VNFs run all DHCP sessions for all vCPEs services and communicate with the vCPE-control services. Communication includes the vCPE-control setting the configuration to the DHCP VNFs (like the subnet and IP pool for the DHCP server) and the DHCP VNFs reporting to the vCPE-control on the DHCP session results. The SDN controller sets rules at the pCPE such that DHCP sessions are forwarded to the DHCPs VNF.

The advantage of this approach is in higher efficiency – centralized DHCP services are more resource efficient than distributed services. Yet, as shown above, orchestration and control are more complex.

6.5 CPE-vCPE Architecture Dependencies

Other architectures also may be considered, some being more centralized than other. The pCPE architecture does not dictate a specific vCPE architecture, as long as some conditions are met:

1. The pCPE-vCPE interface is not affected by the vCPE architecture
2. The SDN protocol is generic enough to set rules for the pCPE supporting all architecture

The vCPE orchestration is flexible enough to instantiate vCPEs elastically as pCPEs come and go. At this time, it appears that the optimal vCPE architecture is defined in 6-3 (Split Control and Data Path). Early feedback from vCPE manufacturers shows that having a vCPE data path with dedicated network functions has great efficiency benefits, yet is reasonably simple to manage and orchestrate.

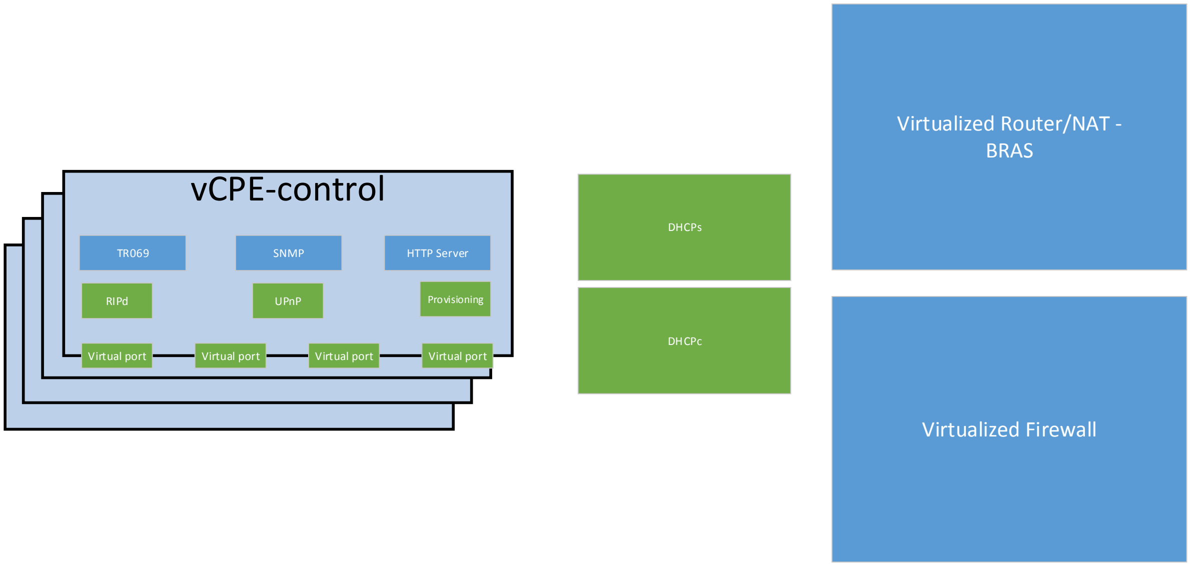

Figure 6-3 vCPEs with control as micro service

7 COMPARISON OF OTT AND SERVICE PROVIDER MANAGED NETWORK FUNCTIONS

With the evolution of broadband service, it is often argued that commercial service provider offerings will transition to being a “dumb pipe.” Network and service virtualization provides advanced capabilities compared to the existing architecture, thus allowing for service providers to bring added value to their subscribers that would help compete with Over-The-Top (OTT) service providers. This section discusses those capabilities and highlights inherent advantages of MSO hosted SDN/NFV network compared to OTT VNFs.

7.1 Personalized Services

Making the pCPE a part of the SDN domain opens the door for user and physical port aware virtualized services. As an example, device aware parental control policy can be performed as a part of a parental control VNF, as the device MAC address and optionally the port information can be carried with the packet and traverse SDN switches.

If the pCPE is not a part of the SDN domain and acts as an Embedded Home Gateway a lot of valuable information is lost, including the user device identification and even IP address (in case of NAT) and the physical port. Over The Top (OTT) providers relying on embedded gateways do not have access to this valuable information and can therefore offer limited services. To follow the parental control example, an OTT parental control VNF can only associate traffic with a CPE device, not with the actual end device or the physical port.

7.2 Responsive Services

OTT also has the major disadvantage of remote data centers. While the MSOs hubs are located both physically and topologically within the vicinity of the subscriber premises, public clouds are located further away and limited in ability to provide services which require fast response time.

7.3 End-to-End Service Assurance

Many times it is desirable to assure the availability of a network service. Best effort services do exist but are harder to monetize and often provide poor user experience.

SDN network flows are set by the SDN controller. Apart from setting the path, the SDN controller can also set up end-to-end QoS on a per-flow basis. The MSO is a fully managed network so service assurance can be provided while in the MSO network using SDN techniques.

OTT services reside on the public Internet and are therefore not subject to controlled QoS and are at the mercy of the Internet bottlenecks and peak loads.

8 TTM AND FEATURES DEPLOYMENT BENEFITS – QUANTITATIVE ANALYSIS

The major benefit called out for vCPE is service agility. The new approach enables a new multi service x86 platform which enables seamless introduction of new services, service upgrades and simplified end user support. By moving to vCPE-enabled services, some analysts [4] believe CSPs can achieve up to ~75-80% of CPE-related cost savings per household per year after the migration of a majority of the customer base. Others [10]claim a cost saving of up to 24%.

The move can be broadly categorized under 3 headings.

1. Flexible Multi Technology CPE

The CPE can be used in both modes, today where it runs all the legacy functions or as described above where some of the CPE functions are moved to the core network. This Home Gateway can now act as L2 or L3+ device allowing the network edge to orchestrate the appropriate level of intelligence and the Home Gateway can provide connectivity functionalities including an access point (WLAN), switching (Ethernet) and modem (xDSL, Cable), which from the OpEx perspective makes it cheaper to resource, manage, and repair.

2. Cloud vCPE solution

The ability to move advanced IP functions (L4-7) from the Home Gateway to to the Provider Edge or Point of presence (PoP), deployed as virtualised network functions (VNFs) running on COTS servers. Example functions include: CGNAT, Firewall, routing, DHCP, Web management, UPnP.

3. OSS evolution and Customer Support

A fully SDN orchestrated CPE enables automation of the vCPE service deployment, troubleshooting and service upgrades. L2 visibility to the home which allows full visibility to the orchestration layer enables CPE configuration, testing and troubleshooting and operations through a customer self-service portal. This reduces the number and duration of customer calls and truck rolls.

There are varying estimates from different analysts and reports on TCO savings based on the headings above, but they assume varying consumer scenarios such as double-play routers and triple-play set-top boxes. One thing that is consistent is that ability to co-locate services and service deployment, reduce CPE boxes that need to be managed, deployed and maintained, and the ability to up sell services (e.g. try before buy and profit share with OTT's) offers ~75-80% total cost savings per household

New Value Added Services

Operators do not necessarily know which new services may arise, but they do know they can no longer operate in the traditional mode of 18-24 month HW development cycles to enable such service introduction. They must provide flexible multi-locational platforms which enable faster Cloud/OTT like service deployment capabilities which enable new service introduction or indeed the flexible platform onto which external platforms can deploy revenue sharing service models.

Roll back

CSP's nominally do 2-3 firmware updates per year. The logistical, technical and financial impact of an update going wrong is such that every upgrade is a heavily controlled and rehearsed event. The cost of a failed upgrade carries a huge opex and capex penalty to the operator resulting from a rise in support calls, truck rolls for repair/replace and the potential for customer loss. This is not the case for cloud services as these can be spun up and down without the end user even noticing. Cloud service providers can update several times a week without the user seeing any noticeable impact to service.

The VNF eco system

We are already seeing a large ecosystem develop in this space. Traditional telecom vendors such as —Alcatel-Lucent, Cisco, Ericsson, Fujitsu, Huawei, Juniper, NEC, and ZTE—remain the favored vendors from which operators will buy SDN and NFV software to satisfy and fulfill the vCPE enablement and roll out.

Ecosystem vendors are likely to deploy multiple functions (such as firewall, IP/MPLS VPNs, load balancing, QoS support, VPN termination, CG-NAT, DPI, IDS/IPS, WAN optimization controller, or BRAS/BNG) from physical edge routers to software vRouters or VNFs running on commercial servers

Taking Nokia/ALU as an example, they are demonstrating six x86 servers behaving as a single router that can offer up to 2Tbit/s throughput, with the control plane and EMS [element management system] regarding it as a single instance. Alcatel-Lucent has architected its routing software to get the best performance, resilience and reliability on general purpose (x86) compute platforms. The company is demonstrating [11] 320G half-duplex, or greater than 2x better than competitor offers, for a virtualized Provider Edge routing application in a single x86 server.

Brocade in conjunction with Telefonica have shown their Vyatta 5600 vRouter on a commercial off-the-shelf based x86 server within a Red Hat KVM environment. Tests have shown 8x the performance in the control plane over competitor offers [12]. Deployed as a single virtual machine, the Vyatta 5600 was able to support all of the server's available ports at line rate by hitting the 80Gbit/s mark, Brocade actually exceeded its own original goal, which was to prove that a software router can support the 10Gbit/s performance that is mainstream in carrier environments

Table 1 – Scope of Visibility per CPE Technology

	L3 NAT decisions	L2 switching decisions	Level of visibility from outside the home
Embedded Home Gateway	By the pCPE	By the pCPE	Traffic between the home and the internet
Virtual Home Gateway	By the vCPE	By the pCPE	Traffic between the a client within the home and the internet
Virtual Home Gateway with selective hair-pinning	By the vCPE	By the vCPE	Traffic between clients and the internet or between themselves

9 CHALLENGES

9.1 Security and Privacy, DOS

Different degrees of virtualizing Home Gateway functions imply different levels of visibility from outside the home. We expect that the Virtual Home Gateway will be tightly linked to new services for subscribers to allow it. These services may include community WIFI and home automation with advanced network protection.

An architecture where the pCPE is a complete slave of the vCPE (as proposed in 4-3) exposes some security and privacy vulnerabilities that have to be addressed. Impersonation of the vCPE and the SDN controller can result in having subscriber traffic, including LAN to LAN traffic, diverted to an arbitrary path chosen by the attacker. It is expected therefore that virtual Home Gateway operation would imply two-way authentication between vCPE and pCPE. New generations of threats could include *man in the middle* attacks between vCPE and pCPE. Also, DDOS attacks by pCPE on the shared L3 NAT resources of the vCPE will drive motivation for authenticating “traffic patterns” generated by pCPE – even when not violating the bandwidth service level.

Another effect of the proposed architecture is the potential exposure of LAN to LAN traffic to the service provider via hair-pinning. Although it does allow the exposure of more content to the service provider, the convention today is that all traffic being sent in the clear is exposed. Specifically, all Internet traffic is exposed to the service provider today so the additional exposure induced by hair-pinning is not likely to be perceived as violating the level of privacy provided today.

9.2 Intelligent Hair-Pinning Under Limited Uplink Resources

The hair-pinning mode of operation between vCPE and pCPE makes it possible for Service Providers to associate advanced processing functions to specific traffic flows between two clients connected to the same residential local area network – wireless and wireline.

Advanced processing of LAN to LAN traffic could be used by holistic smart home network security solutions designed for protection in an IOT environment with no assumption on the source or the capabilities of the “Things.” This includes the scenario of friendly visitor clients being legitimately authorized to onboard one’s home network. Such solutions will rely on consuming just as much traffic as processing and bandwidth allow in search for suspicious traffic patterns across the residential network.

As broadband technologies are asymmetrical and as home traffic volume and traffic peak are increased, managing upstream resources under hair-pinning operational mode between a vCPE and a pCPE becomes a new challenge. We expect that the vCPE would have to be real-time aware to actual utilization of the upstream and to the type of clients involved in each session. As an example, traffic between the Refrigerator and the TV is likely to be more suspicions and might require advanced inspection. The vCPE will therefore be smart enough to make real-time intelligent decision with respect to which LAN to LAN traffic should be processed on the vCPE platform and which should be offloaded for switching by the pCPE.

Designers of next generation broadband technologies should be aware to possible increased demand for improved upstream/downstream ratio.

10 SUMMARY

Virtualization of the home gateway as a broadband service provider has many advantages including significant reduction in OpEx, service agility, enablement of compute intensive services and enablement of services at the subscriber LAN scope. Several network architecture alternatives are made possible, each with its own advantages and complexities. Rather than pointing to an optimal solution, this paper compares the alternatives and analyzes them.

We believe that the transition of the home gateway architecture from an embedded one to a virtualized one will bring many benefits to the broadband service provider. We also believe that the technological infrastructure for the transition is for the most part ready and that the transition should happen in the near future. We encourage Service Providers and ecosystem vendors to contribute to making this transition as soon as possible.

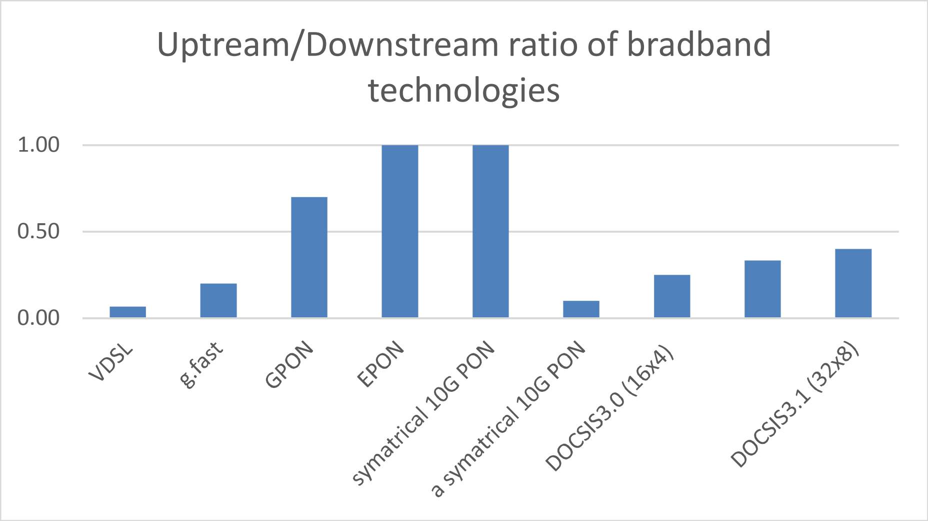

Figure 9-1 - Up/Down Bandwidth Ratios per Technology

11 REFERENCES

[1] Nagesh Nandiraju , Ph.D. and Sebnem Ozer, Ph.D.; Arris Inc., "Applying the Software Defined Networking Paradigm to MSO Commercial networks," in *NCTA Spring Technical Forum*, 2013.

[2] Nagesh Nandiraju Ph.D., Yiu Lee and Jorge Salinger; Comcast Cable, "Delivering Seamless Subscriber Aware Services over Heterogeneous Access Networks using a SDN and NFV framework," in *INTX Spring Technical Forum*, 2015.

[3] Cable Labs, *Data-Over-Cable Service Interface Specifications MAC and Upper Layer Protocols Interface,* Cable Labs, 2015.

[4] G. Ragoonanan and Y. Gorkem, "vCPE services business case: potentially billions of dollars payback for fixed CSPs," Analysys Mason, 2015.

[5] D. Farinacci, T. Li, S. Hanks, D. Meyer and P. Traina, "IETF.org," IETF, March 2000. [Online]. Available: tools.ietf.org/html/rfc2784.

[6] W. community, "Wikipedia," Wikipedia, [Online]. Available: en.wikipedia.org/wiki/Internet_Mix.

[7] "openvswitch.org," [Online]. Available: openvswitch.org.

[8] European Telecommunications Standard Institute, *Network Function Virtualization (NFV); Architectural Framework,* ETSI, 2014.

[9] LinuxContainers.org, "LinuxContainers.org," [Online]. Available: linuxcontainers.org.

[10] "The reality of cost reduction," HP, 2015.

[11] Light Reading, 11 12 2014. [Online]. Available: http://www.lightreading.com/nfv/nfv-elements/alcatel-lucent-joins-virtual-router-race/d/d-id/712004.

[12] Light Reading, 19 8 2014. [Online]. Available: http://www.lightreading.com/nfv/nfv-elements/telefonica-proves-brocade-router-performs-for-nfv/d/d-id/710397.