Software Defined *Access* Networking (SDAN)

Taking NFV and SDN to the "Edge"

WHITE PAPER

Contents

Executive Summary

Extending the benefits of Network Functions Virtualization (NFV) and Software-Defined Networking to the edge of the network, including access segments and subscriber premises, offers enormous gains in flexibility and control of broadband access services. This concept is described here as "Software Defined Access Networking" or SDAN.

SDAN virtualizes access-network control and management functions for broadband access, with the goal of streamlining operations, accelerating service creation, enhancing performance, enabling competition among service providers on a shared infrastructure, and ultimately enhancing customer satisfaction with broadband access.

Although this paper focuses principally on DSL access technology, SDAN includes any type of broadband access: digital subscriber lines (DSL), cable modems, fiber-to-the-x (FTTx), wireless, and so on. Virtualization and control of the in-home network also is within the realm of the SDAN. An integral aspect of SDAN is that unlike traditional SDN, advanced management and control down to the PHY level are highly desirable and included features.

This paper presents three network scenarios representing progressive levels of SDAN:

1. A single entity operates the physical access infrastructure and sells access services enabled by that infrastructure to end-users. This entity is thus a physical infrastructure provider (PIP) and a network operator. By centralizing access network control, diagnostics and analytics, the provider can greatly improve operational efficiencies and expand revenue opportunities. Such practices are already applied on tens of millions of DSLs worldwide.

2. A PIP manages a single physical infrastructure over which multiple virtual network operators (VNOs) offer end-user services. This scenario includes bit-stream access and variants such as virtual unbundled local access (VULA) services. SDAN virtualizes the network equipment to enable VNOs to realize a level of flexibility and a level of efficiency which is the same as owning and directly controlling the equipment, while providing for determinacy in the management and performance of the shared, underlying physical network.

3. PIPs and VNOs have access to data regarding consumer experience and broadband performance, which is made available from platforms monitoring end-user devices, services and customer premises networking technology. This additional data is used to better monitor access services and tune performances to meet end-user needs, which expands the VNOs' abilities to improve customer experience and offer new types of services.

1 A Physical Infrastructure Provider (PIP) differs from the concept of an Infrastructure provider in network function virtualization (NFV) terminology [2][3]. In NFV literature, an infrastructure provider is the provider of the compute, storage and network hardware resources that constitute the execution environment for virtual network functions (VNFs).

Introduction

This section gives a brief introduction to access networks, and to SDN and NFV concepts. It also provides a high-level explanation of the three scenarios for SDAN. For more details on industry drivers that lead to the introduction of the new SDAN concept, see the IEEE Communications Magazine article "Software-Defined Access Networks[1]."

1.1. Access Networks

Internet service is delivered to residential consumers through the access network (also known as the "last-mile" connection). Regardless of the transmission medium, access networks contain three fundamental parts: access nodes on the infrastructure side, residential gateways on the consumer side, and the (typically passive) outside plant that connects the two sides.

In DSL networks, the access node is called a DSL access multiplexer (DSLAM); in fiber networks, it is called optical line terminal (OLT); and in coaxial-cable networks, it is called a cable modem termination system (CMTS). The access nodes are connected to the core network through aggregation networks. The residential gateways are connected to consumer devices through the home network.

In some DSL markets that have regulatory requirements for unbundling, multiple entities are involved in the delivery of Internet services. As shown in Figure 1, the copper cable may be owned by a metallic path provider, the access nodes and the aggregation network may be controlled by an access node provider, and the home and broadband networks may be in the control of a VNO. In most markets, the metallic path provider and the

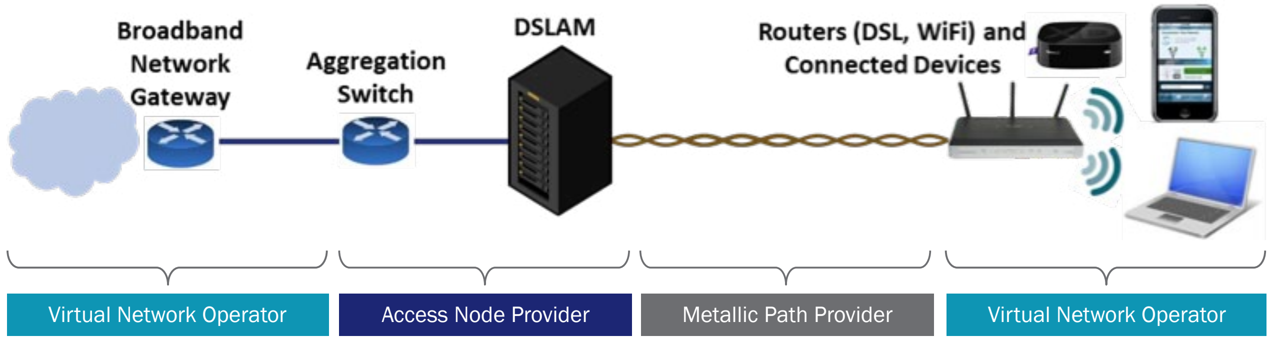

Figure 1: DSL network example with multiple entities controlling different segments

access node provider are the same entity, which this paper defines as the physical infrastructure provider[1] (PIP).

The separation of the PIP from the VNO started with DSL networks, but it is technically possible and is gaining ground for fiber and coaxial-cable networks. Here are some noteworthy related developments:

- Discussions among UK providers for a joint venture for building fiber infrastructure [4].
- Standards proposals for fiber network virtualization [5].
- Research efforts into economics of fiber infrastructure sharing [6].
- Efforts by the Belgian telecommunications regulator to promote competition over the cable infrastructure [7].
- Strategic decision of NBN-Co to use hybrid-fiber-coax (HFC) infrastructure to deliver fast-broadband to Australian retail service providers [8].

1.2. SDN and NFV

SDN is generating a wave of interest as a technology that makes networks programmable and less expensive to build, operate, and upgrade. It is typically described as separating the control and the data forwarding functions of network equipment, and centralizing the control functions of multiple network elements. More generally, SDN is a framework for automatic and dynamic management of multiple network elements. So far, SDN control has referred to functions above the physical layer—such as packet forwarding—and in the core network.

NFV brings many of the virtualization concepts developed for enterprises and data centers into the telecom world, and today is being looked at with growing interest. NFV leverages virtualization technology to consolidate network equipment using commercial-off-the-shelf servers and storage located in data-centers, network centers, and possibly also at user premises. This is expected to lead to reduced equipment costs by reducing requirements for purpose-specific hardware. Furthermore, network and service features defined in software will allow rapid changes in service definitions and sharing of hardware components.

Network functions migrate from embedded firmware in the dedicated hardware of various broadband access network appliances to software running on commodity hardware in a private or public cloud. These virtualized network functions run in a standardized execution environment thus eliminating the dependency between a network function and its dedicated and vendor-specific hardware. In some cases, it is not possible to completely eliminate hardware dependency. This is the case when network functions are only partially virtualized, and may occur when complex signal processing must run locally in real time, or when interfaces between the analog and digital worlds are present. However, even if there are functionalities that are not easily virtualized, it is still possible to eliminate the dependency on vendor-specific hardware by using open interfaces and appropriate abstraction layers for the underlying hardware.

1.3. Three Scenarios for Software-Defined Access Networking

This paper presents three access network scenarios, with progressive degrees of virtualization. These three scenarios correspond to three *possible instantiations* of a SDAN.

The first scenario is that of a network owned by a single operator, where a single entity serves as PIP, i.e. controlling the outside plant and the access hardware, and also as a network operator, i.e. providing services to consumers. This scenario is representative of many current networks in which there is no virtualization of the access network. In this scenario, a management system centrally controls and manages all lines connected to the access hardware. Control functions include the optimization of the physical layer configuration of the broadband connection, for example changing parameters such as data rates and power levels. Management functions include network analytics that drive maintenance operations and marketing campaigns. Centralizing these functions enables the use of advanced algorithms delivering the highest gains, and it also allows for homogeneous, vendor-agnostic management of the access hardware.

Next, a scenario in which multiple VNOs deliver services to end-users over a single infrastructure is considered. The outside plant and the access hardware are controlled by one or more PIPs, who may themselves be VNOs. This scenario applies to current networks where regulatory provisions for competition via local-loop-unbundling (LLU), sub-loop-unbundling (SLU) or bit-stream access exist[2]. This scenario is especially common in regions with facilities-based competition or where Layer 2 competition is mandated by regulators. In this case, a multi-tenant management system enables the VNO to maintain substantial control of services, while providing for determinacy in the operation and performance of the underlying physical network.

Finally, the third scenario is a complete software-defined access network that extends beyond the access segment, and includes the in-home network and possibly application platforms. In this case, a cloud-based system receives data on end-to-end broadband access and application performance. The data is collected by agents in the network, and is used to guide the optimization and diagnostics engines of the broadband network's management system.

2. Single Operator SDAN

This section describes the first SDAN scenario. Typically in this scenario, a single entity serves both as the PIP (controlling the outside plant and the access hardware) and as the VNO providing services to consumers.

2.1. Access Network Control

Access network control is necessary for configuring the access network to support the delivery of broadband services to consumers. Configuration changes are applied to the access nodes, which affect physical layer configuration parameters such as data rates, transmitted power and spectrum, coding schemes, resilience to noise, and latency. For DSL, a predefined set of such configuration parameters is called a *profile*.

2.1.1. Traditional Access Network Control

Traditional access network control relies on the use of default profiles. These profiles, also known as *golden profiles*, have a one to one correspondence with a service product, and they are applied to all newly provisioned lines with a given service product. However, there is vast diversity in the transmission characteristics of lines in the access network, either because of the noise environment (e.g. interference in the home and outside plant), or because of the transmission medium (e.g. variation in the quality of the channel, differences in line lengths, bad contacts or splices, and imperfections of the copper or fiber). The application of a golden profile to such a mixed population of lines results in poor performance for a substantial portion of these lines.

The traditional remedy to the limitations of using golden profiles is the introduction of *maintenance profiles*. Such maintenance profiles are applied to lines that are found to suffer from poor performance. This typically occurs in reaction to customer complaints, or as technicians work to correct issues in the field. This means that a maintenance profile is applied only after the customer has experienced a poor service, and after the provider has incurred significant customer support expenses.

Finally, the manual application of maintenance profiles greatly complicates the enforcement of a network-wide profile policy. Contact center agents and technicians can be inconsistent in following established guidelines, or follow their own preferred practices. As a result, maintenance profiles are improperly or inconsistently used in the network. This weakens the enforcement of a network-wide profile policy, and results in poor customer experience and higher operational costs.

2.1.2. SDAN Control

The modern approach for access network control includes the use of *software-based profile optimization*. Each service product is assigned a set of allowed

2 In the case of bit-stream access, VNOs are often called resellers.

profiles[3], and software applies to every line in the network the profile for that line's service product that best matches the line's specific transmission characteristics. The software algorithms performing this profile optimization rely on historical and current data about the line, and account for the line's service product. The profile optimization process is performed for newly provisioned lines, and for lines that are performing poorly or below their full potential. Mutually interfering lines are also optimized for improved coexistence[4]. In this way, the access network is automatically and continuously adapting to evolving conditions without requiring manual intervention. Additionally, profile optimization can be invoked in real time by contact center agents or by technicians. The principle of profile optimization is illustrated on the right-hand side of Figure 2.

The above methodology ensures that all profile changes in the network follow a set of rules established by the operator and are embedded in the configuration of the profile optimization software. As a result, software-based profile optimization enables *centralized policy management* for access. The operator controls how resources are allocated to obtain the desired outcomes in terms of service reliability, throughput rates, power consumption, and network crosstalk (specifically for DSL). The operator has freedom to define a diverse family of service products (thus increasing consumer choice), and to apply on each line the optimal policy for each such product.

This practice of SDAN control has strong similarities with the fundamental principles of SDN. SDN separates the data plane from the control plane for routing and switching equipment. The control plane is then logically centralized and made programmable. Applications interface only with the control plane, and are not concerned with the forwarding elements that comprise the data plane. Similarly, SDAN control separates the data transmission layer from the control layer. The control layer contains the logically centralized software that implements the operator's policy. The policy is

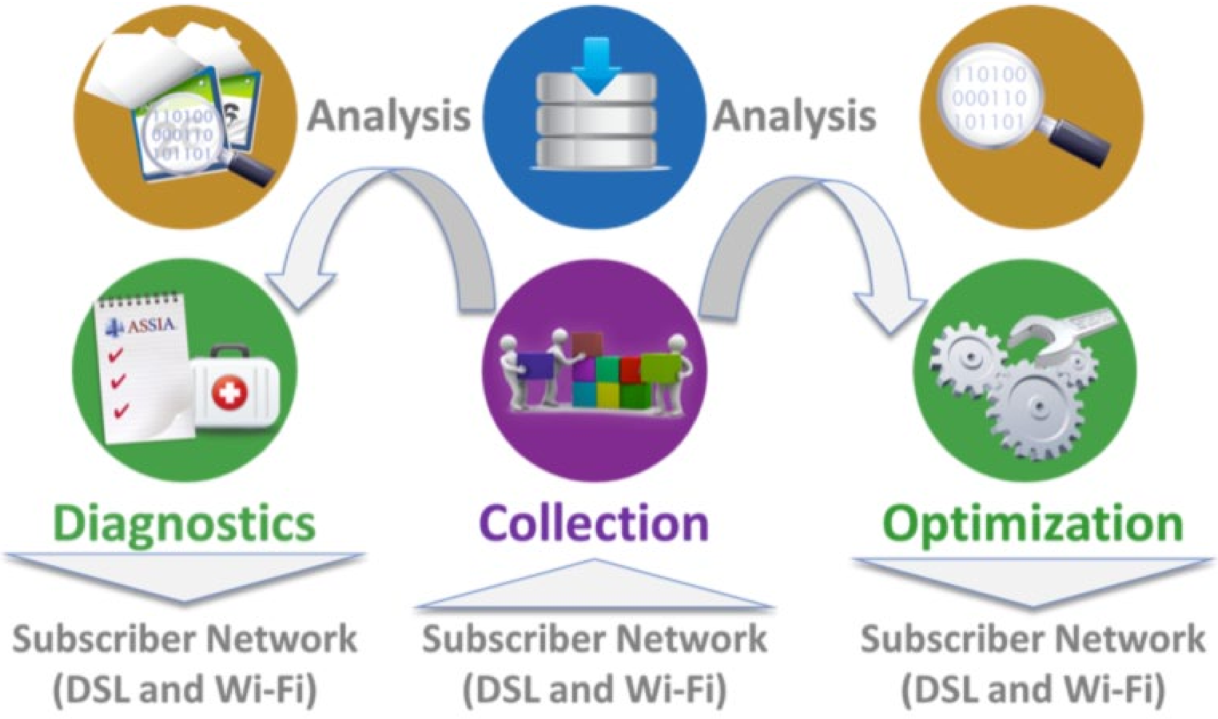

Figure 2: Software-defined Diagnostics (left-hand side) and Profile Optimization (right-hand side)

3 Older equipment limits the number of profiles that can be stored to as few as 40. Newer equipment allows much larger numbers of profiles, and allows the definition of sub-profiles for even higher flexibility.

4 The process of joint optimization of profiles of multiple lines is called vector profile optimization. It applies to mutually interfering DSLs, but also to shared media such as GPON and coaxial cable access systems.

applied homogeneously to the entire access network, which can include any mix of access equipment from multiple vendors.

2.2. Access Network Diagnostics

Access network diagnostics is the analysis of data collected from the access network for resolving service issues, for qualifying customers for services, and for planning network maintenance and upgrades. Data are collected from the access nodes, and may additionally be collected from test equipment connected to the outside plant (for example, specialized hardware for DSL metallic testing, or for optical time-domain reflectometry (OTDR)).

2.2.1. Traditional Access Network Diagnostics

Traditional access network diagnostics rely heavily on hardware systems, for example on specialized test heads for DSL metallic testing, or on handheld devices carried by technicians. Such hardware systems have high costs for deploying and for operating, and often require replacement or expensive upgrade programs for supporting newer generations of access technology. They are used for performing invasive tests, which typically cause service disruption, and must therefore include the prior consent of the affected customer.

More recently, access equipment manufacturers have built access nodes that integrate data transmission functionality with testing capabilities (e.g. metallic loop testing (MELT), and single-ended loop testing (SELT), in DSL; or OTDR in GPON). But the results delivered by such equipment are still physical media diagnostics with limitations similar to those of specialized test hardware.

The use of isolated diagnostics data from specialized test hardware or from access nodes is problematic in several ways:

a. Data from a test performed at a given time may fail to identify intermittent issues.

b. Results from one type of test cannot be combined with other results to improve accuracy and confidence.

c. Results cannot take into account the “context” of the connection, for example, whether a line fault is performance-affecting for the current service product, such as Internet access or IPTV.

In summary, traditional access network diagnostics can be characterized as a *static network function*: They heavily rely on hardware, they are difficult to enhance, and they are isolated from other network functions.

2.2.2. SDAN Diagnostics and Recommendations

The modern alternative to traditional access network diagnostics is the migration of analysis functionality to server infrastructure, while keeping only the necessary raw test-data capture functionality on the access node (or on specialized test hardware). The advantages of this approach are explained below.

SDAN-based diagnostics enable the combined use of raw test-data from multiple sources. For instance, data can be combined from sources such as access nodes, service management platforms, and specialized hardware. Also, different types of test data can be combined, e.g. SELT, MELT and Double-Ended Loop Testing (DELT) data in DSL. Finally, the analysis can take into account historical data to better analyze time-varying effects. The end results are improved accuracy for diagnostics and improved identification of the root causes of degradation.

A software-based framework allows for the expansion of diagnostics functionality for further inclusion in recommendation functions. Recommendations may include the identification of upsell opportunities, based on analysis of the following:

a. the current service of a consumer,

b. the capacity of the line for a higher-tier service, and

c. the historical broadband usage of the connection.

Other types of recommendations or analytics can include the identification of customers with high risk for churn, or the prioritization of certain areas for proactive outside plant maintenance. All of the above functions benefit from network-wide learning, and can easily be adapted to the needs of each operator. The principle of SDAN-based diagnostics and recommendations is illustrated on the left-hand side of Figure 2.

SDAN-based diagnostics and recommendations are a *virtualized network function* (VNF). They are implemented on commodity server infrastructure, they rely on an abstraction layer for the underlying testing hardware, and they can be flexibly defined and updated to meet the evolving needs of operators.

3. Multiple Operator SDAN

This section discusses the scenario in which a network is used by multiple operators. In this scenario, there are at least two VNOs that are providing services to consumers. The outside plant and the access hardware are controlled by an entity, the PIP, which can either be one of the VNOs or a third party.

3.1. Access Network Sharing

Access network sharing is imposed by regulation in a number of broadband markets, e.g., LLU, SLU and bit-stream access, or it can be the result of voluntary operator collaboration. Subsection 3.1.1 describes the principles of traditional access network sharing, while subsection 3.1.2 explains how the individual needs of operators for configuration, diagnostics and remediation of connections in the access network can be supported for each VNO using a multi-tenant management platform.

3.1.1. Traditional Access Network Sharing

Traditional access network sharing often relies on regulations for LLU. In the vast majority of markets, physical wire connections between local exchanges and customer premises are owned by incumbent providers. LLU regulations require that competing providers be granted access to these connections. The competing providers lease the wire connections from the incumbents to deliver services to their customers and install their own access hardware.

In recent years, copper wire connections have been shortened by installing fiber connections between local exchanges and *cabinets* in the outside plant. The shortening of the wires enables newer access technologies that deliver higher-speed services. In such cases, SLU regulations require that competing providers be granted access to the sub-sections of wires between cabinets and customer premises. The competing providers must again lease the wire sub-section and install their own access hardware at the cabinet location.

Both LLU and SLU have been applied to DSL networks. In LLU and SLU, competitive operators deploy their own equipment in a central office or cabinet, and regulations specify rules for the physical layer characteristics of each provider's operations on the wire connections. As an example, transmission power and spectra rules are imposed to minimize performance-limiting interference among wires. Such rules are traditionally designed for the worst case scenario that can be encountered in the network with the assumption that there can be no coordination among operators' uses of the lines. Consequently, access network performance may not be optimum. As discussed in sections 3.1.2 and 3.2.2, performance can be optimized by using a multi-tenant management platform to virtualize the networking equipment.

Traditional access network sharing is applied through *bit-stream access* in some jurisdictions. This is a form of virtual unbundling, where the competing providers are granted access not to the wire connections but to Layer 2 connections between the competing providers' network hardware and the customer premises equipment (CPE). Bit-stream access mainly applies to DSL networks at this time, but its principles can also be used for fiber-based or for coax-based access.

Bit-stream access is often advertised as preferable for deployments of vectored VDSL. This is because the gains from using vectored VDSL rely on coordination among the ports of all equipment at a given cabinet. However, competing providers have expressed concerns that bit-stream access limits their ability to offer differentiated access services. Some have claimed that this approach essentially re-monopolizes the access network and returns complete control to the incumbent provider, thereby reducing meaningful competition[5]. The Broadband Forum has recently released two documents that present a broad industry consensus on best practices for mitigating the effects of un-cancelled crosstalk on vectored VDSL2, so that facilities-based competition and vectoring performance gains can be realized in vectored environments [9][10].

3.1.2. SDAN Network Sharing

There are cases where operators have an incentive for implementing access network sharing in a way that goes beyond the LLU/SLU/bit-stream access models. This is the case when certain broadband infrastructure investments require a high consumer take rate for services provided over that infrastructure to make the investments economically justifiable.

In this case, the physical network is partitioned into virtual networks corresponding to the customers of each of the VNOs. Virtualization provides an abstraction mapping the physical access hardware (e.g.

5 The practice of allowing only one provider deploy a vectored DSLAM at a given site has been described as the "Highlander (There Can Only Be One)" approach.

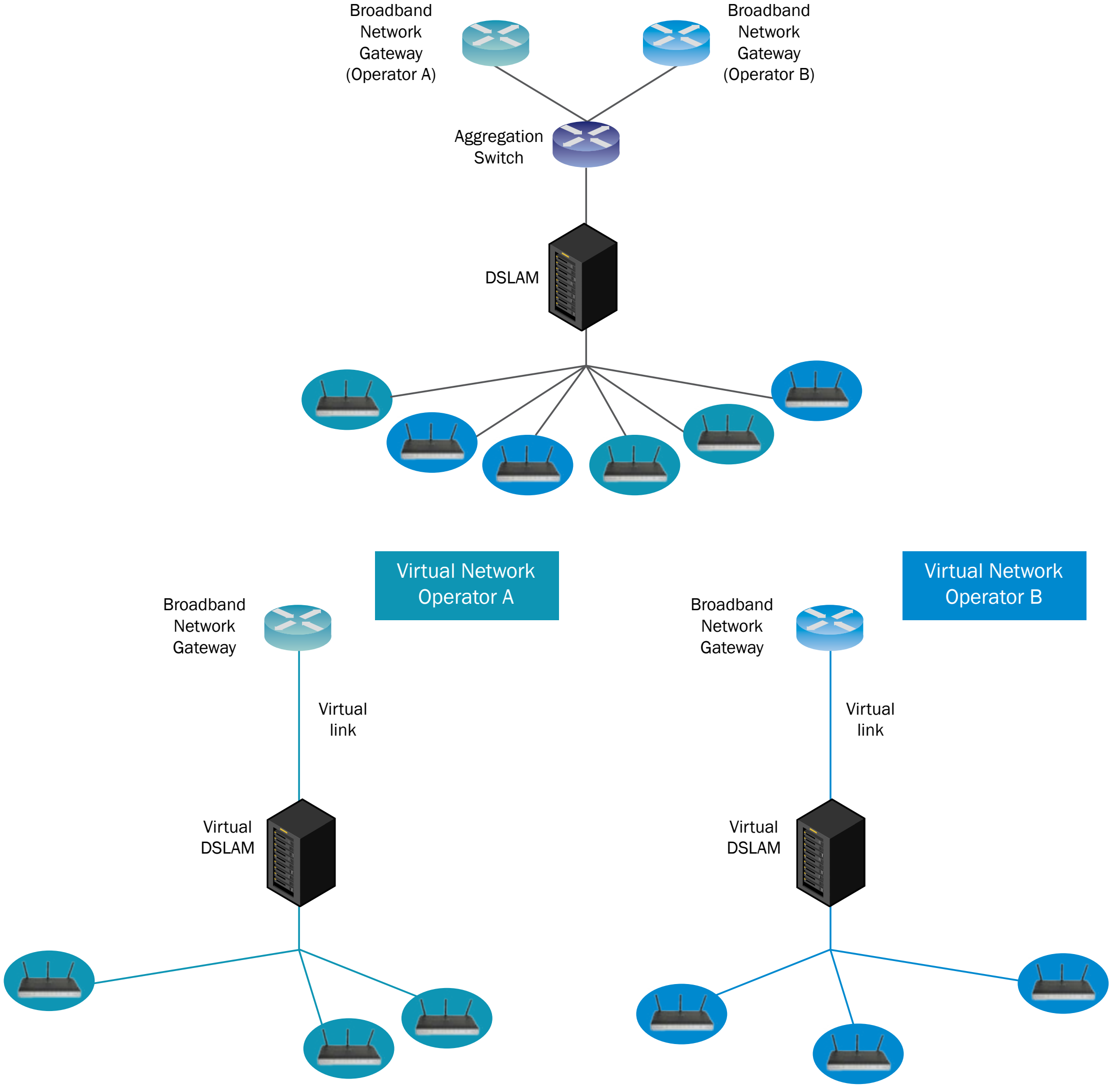

Figure 3: Mapping of Physical Access Network to Virtual Access Networks

DSLAMs or optical network terminals) to virtual hardware. The physical access hardware is mapped to virtual hardware as shown in Figure 3 for the case of DSL, while the virtual hardware can be controlled and managed by each VNO.

The great advantage of this approach is that the virtual hardware can be controlled and managed by the VNOs in a way that provides for determinacy in the operation and performance of the underlying physical network. As explained further in subsection 3.2.2, the VNO has information about and control over the access network, almost as though the VNO owned its own hardware. As opposed to the use of bit-stream access, this model lets the VNO design its own competitive service products and make real-time changes on the network. Variants of bit-stream access, such as virtual unbundled local access (VULA) aim to deliver services "*with a degree of control that is similar to that achieved when taking over the physical line to the customer*" [11]. However, existing implementations of VULA fall short of such promise.

Finally, enabling VNOs to share with each other some information about their own network, global optimization across lines served by different VNOs becomes possible. Such a capability can greatly improve network performance because the multi-user nature of most access channels can make any form of independent management of lines sub-optimal. When all lines are vectored, the channels become decoupled. In this case, VNOs can independently manage lines.

3.2. Configuration, Diagnostics and Remediation Functions for Shared Access Plants

Every provider managing access lines must have processes for configuring these lines to support the services delivered, for knowing the line performance and diagnosing issues, and for responding to remedy issues. Performing these functions within a shared access infrastructure presents additional challenges. The traditional and SDAN approaches for implementing these functions are described below. VULA is assumed for traditional access network sharing, given that it is most often presented as appropriate for next-generation access technologies such as vectored VDSL and G.fast.

3.2.1. Traditional Functions for Shared Access Plants

In practice, VULA gives competing providers very little control over line configuration. Typically, the provider can choose from a limited menu of services for the access segment. These services are characterized by a data-rate range, but have little or no differentiation in terms of quality of service. Data collection schedules, quality of service definitions, service qualification rules, and profile optimization rules are not within the control of the competing provider.

Competing providers also have limited visibility into access network performance. Although some VULA implementations include functionality for the provider to retrieve a pre-defined set of performance quantities and test parameters for a given line, the typical implementation does not include functions to enable the provider to independently diagnose the type or the location of a fault. Providers often resort to their collection of data from their customers' CPEs—and even that may not be possible—to quantify the speed and the quality of the connections. This lack of visibility on behalf of the competing providers often leads to miscommunication, waste, and even conflict, when coordination with the PIP is required to address customer complaints. Also, competing providers are not able to share information to assist in diagnosing the source of performance issues (e.g. identifying a copper cable integrity problem that causes degradation on multiple DSLs used by different providers).

Finally, VULA greatly restricts providers in terms of actions that they can take to remedy access problems. The competing providers are entirely dependent on the PIP for correcting any access problem, and have no way to react or give information to their customers in real time.

3.2.2. Multi-Tenant Platform for Shared Access Plants

Access network virtualization enables the definition of a customized set of VNFs for each of the competing providers operating on a shared physical infrastructure.

Figure 4 shows a multi-tenant platform that provides VNOs with access network virtualization and with management VNFs. The configuration of such VNFs meets the network service needs of each VNO and is specifically authorized to manage aspects of the underlying physical infrastructure. The PIP manages its equipment and authorizes VNO activities to satisfy SLAs in place with the VNOs.

Figure 5 shows another instantiation of a multi-tenant management system, which differs from Figure 4 in that management functions are independently implemented by the VNOs. However, for both models data can be shared among operators, providing more degrees of freedom in optimizing the performance of the physical links than the LLU/SLU/VULA models of sharing the network.

The multi-tenant platform ensures appropriate authentication and authorization processes are in place so that the VNOs always act within their authority (for example, VNO "A" must not be able to know the service products purchased by customers of VNO "B"). Also, data sharing and coordination among the VNFs used by each VNO are supported. This is done through communication between functions of the multi-tenant management system in the case of Figure 4, and through communication interfaces defined at the management and OSS layers in the case of Figure 5. Flexibility in composing VNFs is essential to VNOs, as some of them may wish to leverage investments they have already made in service provisioning and management tools. In this case, the network data made available to VNOs by the multi-tenant platform can aid those provisioning and management tools in a way that would not be possible under LLU, SLU, or VULA regimes.

The multi-tenant platform gives VNOs a wide range of capabilities for configuring their respective virtual access networks. A VNO can define service products along any desired dimension, such as speed, latency, stability, and quality of service. It can choose the schedules for

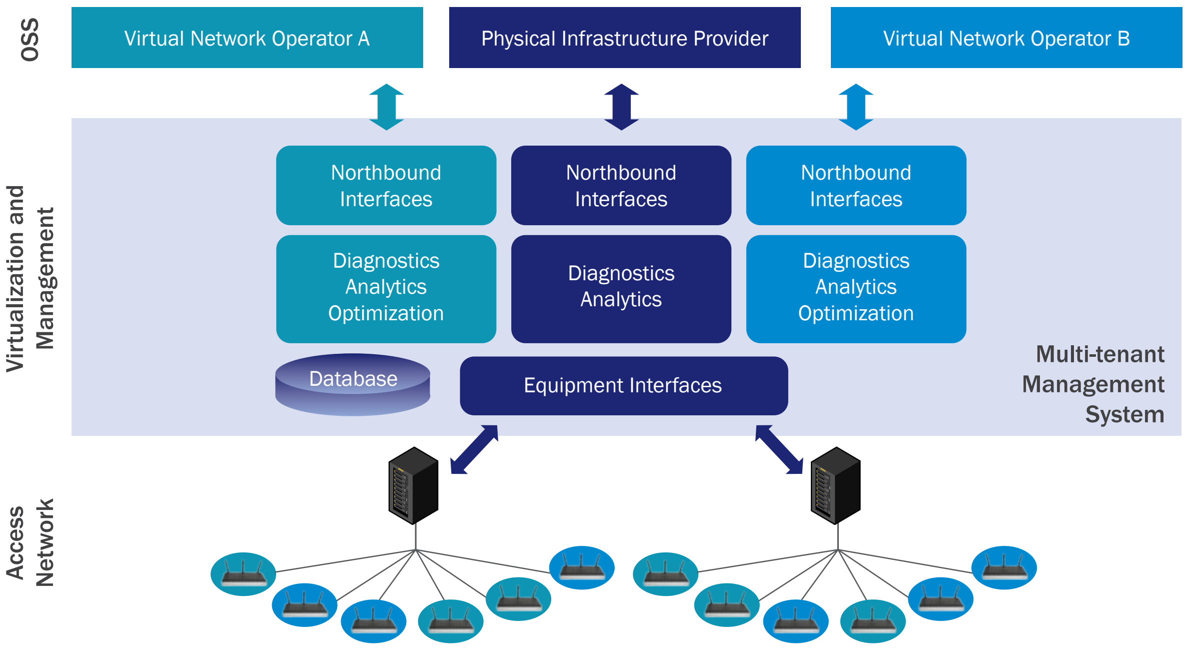

Figure 4: Management functions are part of the Multi-tenant Management System

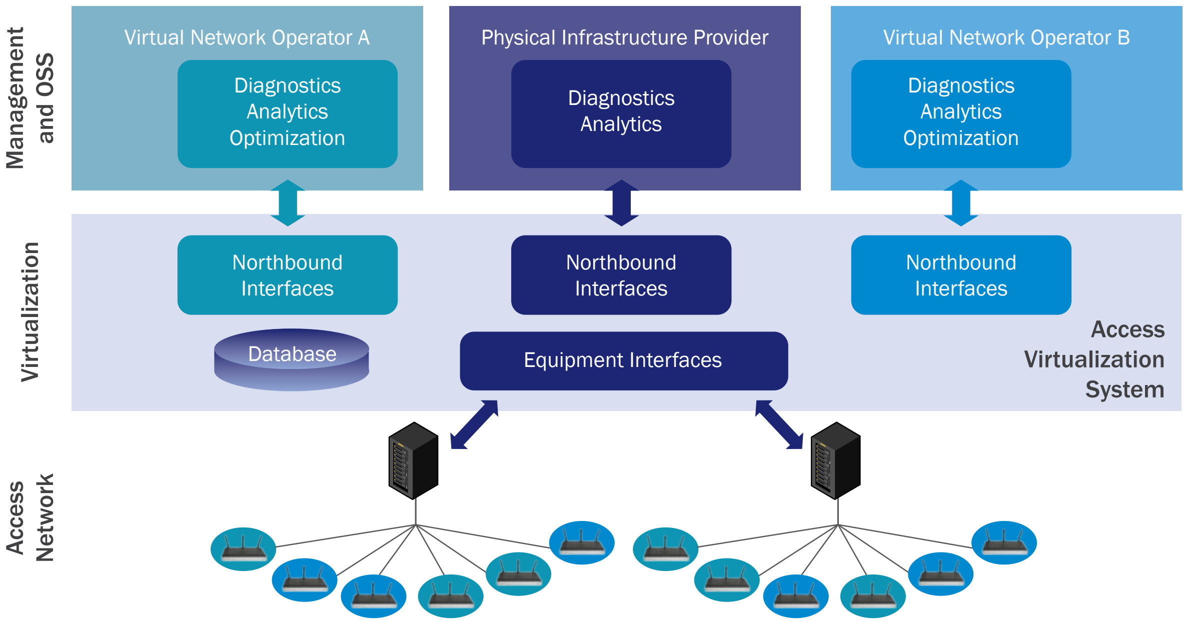

Figure 5: Management functions are outside the Access Virtualization System

proactive data collection and analysis of the network, and can apply its own service qualification criteria. It can choose the rules for profile optimization, and can control the set of profiles to use. Choices with respect to profiles are made within an envelope agreed by the VNOs and the PIP. For example, providers can agree on rules for power and spectrum usage on a per-line or a per-node basis.

VNOs are also given capabilities for viewing and analyzing their virtual access networks. A VNO can retrieve any type of DELT, SELT, or MELT result, together with the diagnostics analysis for a line that it controls. Diagnostics analysis can include localization information, for example combining neighborhood data in a way that hides proprietary information about lines of other providers while improving diagnostic accuracy. Such analysis can identify lines of a VNO that cause harm to lines of another VNO (in such cases, automatic optimization may act to minimize any disruption). Furthermore, a VNO can obtain service recommendations, and view the history of past profile optimizations, as well as derive network-wide analytics, helpful for tasks such as identifying customers that are at risk of churning, or customers that are good candidates for service upsell.

The access network virtualization provided by the multi-tenant platform enables VNOs to act in real time in ways that are not typically possible with traditional VULA. A VNO can take an action such as resetting a port (e.g. when a DSL is stuck at a low rate), or forcing a maintenance profile (e.g. while a technician performs line testing). A VNO can request a real-time data collection and a corresponding diagnostics analysis to check on the latest status of a line. Going a step further, a VNO can obtain recommendations for use by a contact center (for example, to make a decision on dispatching a technician) or by the field force (for example, to verify that the applied fix has corrected a prior issue). Finally, a VNO can initiate a real-time profile optimization. Such optimization can correct a performance problem without the need for a dispatch, or can restore optimum line speed shortly after a technician has applied a fix.

4. SDAN to the Home

4.1. Traditional Management of the Broadband End-User Experience

Management of broadband access services has traditionally extended to the modem at the subscriber's premises, but not beyond. In certain scenarios, notably IPTV in the U.S., service providers also have managed those portions of the premises network assigned to delivering the video signal. But for the vast majority of broadband connections, the network operators have limited visibility into the performance of the home network. This may have seemed adequate at the time when there were only one or two personal computers at home connected via Ethernet cable, but is extremely problematic with the explosion of home devices connected to the internet via home Wi-Fi. Network operators are reporting that customer service calls related to Wi-Fi problems have already reached 50 percent of the total. Cisco predicts that "By 2017, there will be 5 devices per connection for every internet user" [12]. Such device density will necessitate management of the home network for monitoring and improving the quality of the broadband service.

Additionally, traditional management of broadband does not include monitoring of the customer's broadband experience with over-the-top (OTT) services. Typical "speed-test" throughput measurements show very poor correlation with actual quality of service. Network operators have a poor understanding of end-users' broadband experience, and their ability to adapt services in response to such experience is typically limited. Sandvine's Global Internet Phenomena report shows that real-time entertainment already represents 67.4 percent of peak period downstream traffic for North American fixed access networks [13] (see Figure 6). The same report notes that Netflix and YouTube now account for 50.31 percent of all downstream traffic. The growth of such OTT services points to the need for improved management of the broadband end-user experience.

4.2. Improved Management of the Broadband End-User Experience with SDAN

The SDAN provides a natural framework for fusion of broadband service performance data, both *horizontally* and *vertically* across the access and service networks. Figure 7 depicts the architecture for leveraging end-user network diagnostics data. It exposes performance and diagnostics data from multiple network elements and services of the end-user to broadband access management platforms via a secure and published API.

In the horizontal direction, data from agents in the home (but possibly also from the core and edge) portions of the network are combined to assess performance and identify underperforming segments. For example, throughput bottlenecks can be identified by comparing speed and oversubscription bottlenecks along the path to the end-user.

Throughput, connectivity and latency from an access point to a Wi-Fi-enabled device can be measured accurately and continuously. Wi-Fi coverage and interference problems (causes of repeated user frustration) can be identified, and actions for remediation can be recommended. Wi-Fi performance data

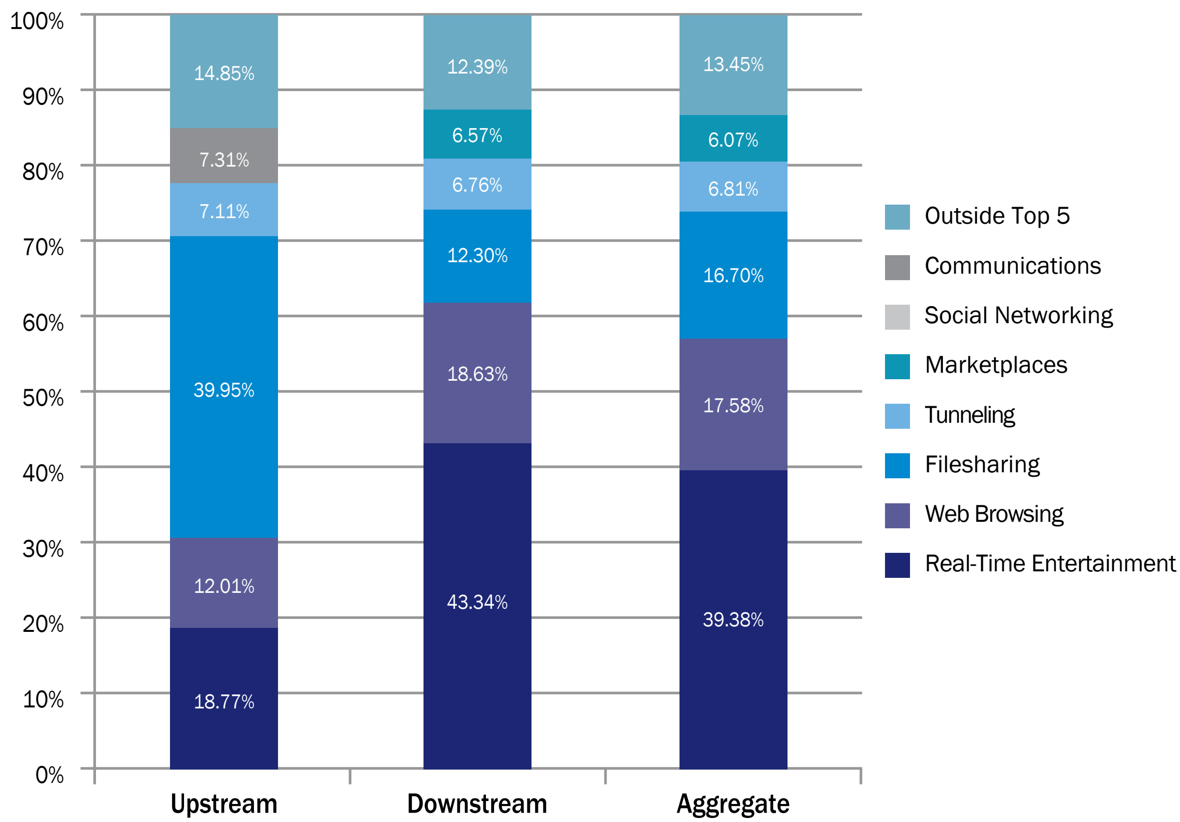

Figure 6: Peak Period Aggregate Traffic Composition - North America, Fixed Access (Sandvine, 2H 2013)

from multiple access points in a multi-dwelling unit can be used to identify radio-frequency or other Wi-Fi interference that degrades Wi-Fi throughput. Business intelligence reports can be generated for network operators, for example to monitor Wi-Fi congestion levels over time, over device-type, or over service product.

In the vertical direction, data from agents in devices at different network layers are compared to assess user experience and develop strategies for optimization. For example, excessive MPEG decoder errors combined with an absence of congestion indications indicate that a higher level of error protection is required on one or more constituent links. The SDAN is particularly helpful in fusing broadband service performance data if multiple service providers are involved (e.g. Verizon for broadband service and Netflix for movie streaming). In fact, historically, end-users with multiple service providers have been unhappily responsible for themselves coordinating repair and troubleshooting among their providers. The SDAN can be employed to create an alternative and more efficient approach where software agents in the home report back to a central database the performance data collected across networking devices owned or controlled by multiple service providers (e.g. a broadband gateway for broadband service and a streaming box for entertainment content). The multi-tenant management system can then set forth actionable recommendations for allowing VNOs to control and fine-tune the access network.

Data fusion is particularly useful when CPEs and Wi-Fi access points are deployed by different service providers as they enable fault correlation to uncover an underlying and common cause for an issue affecting multiple end-users. This is particularly useful in scenarios of multiple users sharing common resources (e.g. interfering Wi-Fi access points, congested coaxial-cable or GPON networks, DSL pairs with crosstalk) where remediation requires simultaneous actions to be taken on multiple devices.

5. Recommendations for Network Operators

This white paper illustrates the concept of *Software Defined Access Networking (SDAN)*, and analyzes three typical scenarios.

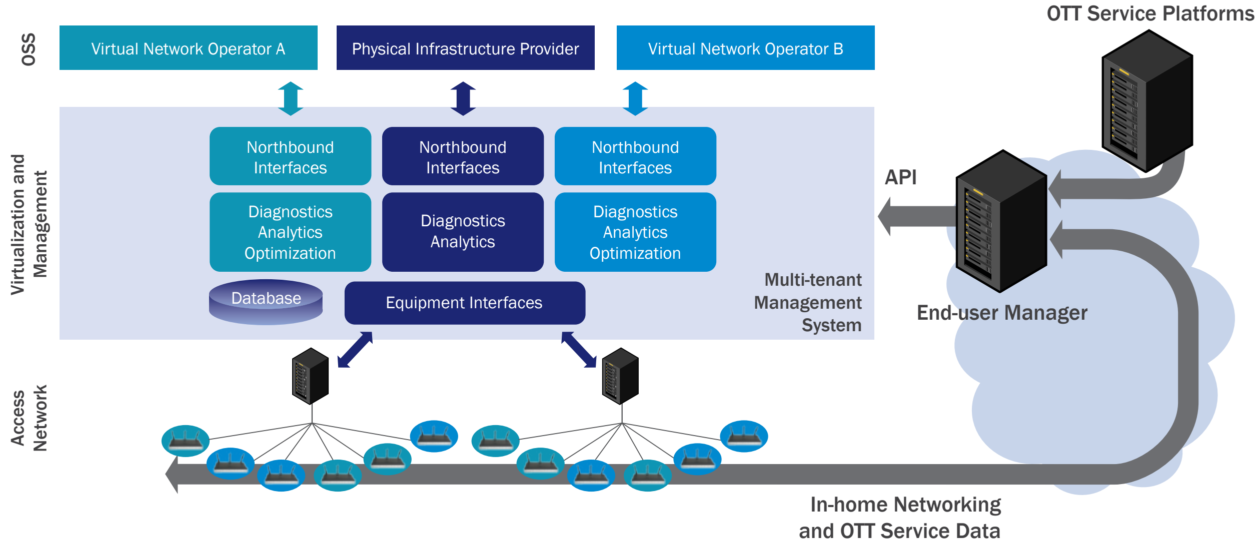

Figure 7: End-user network management server with APIs interfacing with the broadband access management server

The first scenario corresponds to broadband access network management, where a single entity is both the PIP and the network operator. Best practices for such management include the use of software-based optimization for centralized control, and the use of a software-based framework for producing network diagnostics and recommendations. These practices are already applied on tens of millions of broadband lines worldwide. It is recommended to network operators that have not yet adopted such practices to make their selection of a commercially available system based on the following fundamental criteria:

- Level of multi-vendor support for hardware equipment.
- Length and depth of experience with broadband line management.
- Capabilities for automated and fast optimization.
- Proven ability in translating low-level line data into network recommendations and analytics.

The second scenario corresponds to broadband access network management, where multiple VNOs share a single physical infrastructure. The adoption of a multi-tenant management platform empowers VNOs to deliver innovative service products with much greater flexibility than today's bit-stream access. At the same time, the advanced diagnostics and recommendation capabilities that are made available to these operators reduce cost pressures on the PIP. Operators sharing a single physical infrastructure should demand a software-based virtualization layer that gives them network access similar to what would be possible if they owned the access hardware. Infrastructure providers should plan for installing such a layer, so that they can pass more of the network responsibility to the operators.

The third scenario applies to an evolved broadband access management, where the network operator requires management to extend into the home network and to include end-user customer experience with OTT services. Every network operator should have a plan in place for moving towards this newer management paradigm. For creating an agent infrastructure for data collection from network devices, network operators should choose partners with the following strengths:

- Understanding of the needs and requirements of both service providers and end-users
- Management expertise spanning both the access and home networks
- Hardware-vendor independence

For more information on products for SDAN, please visit the following web-pages:

http://www.assia-inc.com/

http://www.assia-inc.com/about-us/contact-us/

References

[1] K. Kerpez, J. Cioffi, G. Ginis, M. Goldburg, S. Galli, and P. Silverman, "Software-Defined Access Networks," *IEEE Communications Magazine*, vol. 52, no. 9, Sep. 2014; retrieved from: http://www.assia-inc.com/technology/knowledge-center/#white-papers

[2] *Network Function Virtualisation: Architectural Framework*, Group Specification, ETSI GS NFV 002, October 2013

[3] *Network Function Virtualisation: Terminology for Main Concepts in NFV*, Group Specification, ETSI GS NFV 003, October 2013

[4] "BSkyB, TalkTalk, CityFibre joint venture challenges BT's broadband hold", FierceTelecom, April 16th, 2014

[5] Bruno Cornaglia, Vodafone Group Services, "Fixed Access Network Virtualization", Broadband Forum Contribution bbf2014.084.00, February 14th, 2014

[6] Juan Rendon Schneir, Yupeng Xiong, "Cost Analysis of Network Sharing in FTTH/PONs," IEEE Communications Magazine, August 2014

[7] "Regulating the Belgian cable," Reinhard Laroy, BIPT, TNO DSL Seminar, June 16th, 2014

[8] *National Broadband Network Strategic Review*, NBN-Co, December 2013; retrieved from: http://www.nbnco.com.au/content/dam/nbnco/documents/NBN-Co-Strategic-Review-Report.pdf

[9] *An Overview of G.993.5 Vectoring – Issue 2*, Marketing Report, Broadband Forum MR-257i2, 2014

[10] *Techniques to Mitigate Uncancelled Crosstalk on Vectored VDSL2 Lines*, Technical Report, Broadband Forum TR-320, 2014

[11] *Review of the Wholesale Local Access Market – Statement*, Ofcom Statement, 7 July 2010; retrieved from: http://stakeholders.ofcom.org.uk/consultations/wla/statement

[12] Cisco VNI Global IP Traffic Forecast, 2012-2017 (2013)

[13] Global Internet Phenomena Report, Sandvine, 2nd half 2013

Acronyms and Abbreviations

API	Application Programming Interface
CDN	Content Distribution Network
CMTS	Cable Modem Termination System
COTS	Commercial Off-The-Shelf
CPE	Customer Premises Equipment
DELT	Double-Ended Line Testing
DSL	Digital Subscriber Line
DSLAM	DSL Access Multiplexer
FTTH	Fiber-To-The-Home
FTTN	Fiber-To-The-Node
FTTx	Fiber-To-The-x
G.fast	ITU-T Recommendation G.9701, Fast Access to Subscriber Terminals
GPON	Gigabit-capable Passive Optical Network
IPTV	Internet-protocol Television
LLU	Local-Loop Unbundling
MELT	Metallic Line Testing
MPEG	Moving Picture Experts Group
NFV	Network Function Virtualization
OLT	Optical Line Terminal
ONT	Optical Network Terminal
OSS	Operations Support System
OTDR	Optical Time-Domain Reflectometry
OTT	Over-The-Top
PIP	Physical Infrastructure Provider
QoS	Quality of Service
SDN	Software-Defined Networking
SDAN	Software-Defined Access Networking
SELT	Single-Ended Line Testing
SLU	Sub-Loop Unbundling
VDSL	Very-high-speed Digital Subscriber Line
VNF	Virtualized Network Function
VNO	Virtual Network Operator
VULA	Virtual Unbundled Local Access
WiFi	802.11 family of IEEE standards

United States
333 Twin Dolphin Drive, Redwood City, CA 94065
Tel: 1-650-654-3400 • Fax: 1-650-654-3404

Europe
Calle Maria Tubau 3, Madrid 28050, Spain
Tel: +34 914842940 • Fax: +34 913446182

China
Suite 470, F/4, Beijing Sunflower Tower No.37, Maizidian Street
Chaoyang District, Beijing 100125, P. R.China
Tel: +86 10 85276788 • Fax: +86 10 85276488

Rev 092514 P/N MC-W1-0115-01-00